Co-efficiency

An empirical investigation of the principle of rational joint action

Georgina Török

Submitted to
Central European University
Department of Cognitive Science

In partial fulfillment of the requirements for the degree of Doctor of Philosophy in Cognitive Science

Primary supervisor: Natalie Sebanz

Secondary supervisor: Gergely Csibra

Budapest, Hungary

2021

Declaration of Authorship

I hereby declare that this submission is my own work and to the best of my knowledge it contains no materials previously published or written by another person, or which have been accepted for the award of any other degree or diploma at Central European University or any other educational institution, except where due acknowledgement is made in the form of bibliographical reference.

The present thesis includes work that appears in the following articles:

<u>An updated version of Chapter 2 has been published as:</u>

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

Author Contributions:

G. Török developed the study concept, designed the task, and collected, analyzed, and interpreted the data under the supervision of N. Sebanz and G. Csibra. G. Török, B. Pomiechowska, and G. Csibra wrote the scripts for the experimental task. All the authors worked on the manuscript and approved the final manuscript for submission.

<u>Chapter 3 has been submitted for publication:</u>

Török, G., Stanciu, O., Sebanz, N., Csibra, G. (under review). Computing joint action costs: Co-actors minimize the aggregate individual costs in an action sequence.

Author Contributions:

G. Török and O. Stanciu developed the study concept, designed the task, analyzed, and interpreted the data under the supervision of N. Sebanz and G. Csibra. G. Török and G. Csibra wrote the scripts for the experimental task. G. Török collected the data. All the authors worked on the manuscript and approved the final manuscript for submission.

Georgina Török

Abstract

People routinely engage in joint actions, coordinating their actions with others to bring about change in the environment. Previous research found that actors plan joint actions with their partner in mind: they take over some of the effort that the partner would have to expend to complete an action sequence. At first glance, this simply reduces the individual action costs of the interaction partner. In this dissertation, I examine the alternative explanation that this kind of behavior is part of rational decision-making in joint action planning. I propose the *principle of rational joint action* as the normative principle underlying action planning: according to this, joint actions should be performed in a *co-efficient* way, by minimizing the joint costs of the action sequence. This would make cooperation more instrumentally efficient and predictable for interaction partners. I present four series of behavioral experiments tackling co-efficiency from different perspectives.

In work reported in Chapter 2, I investigated how people distribute the costs of a joint action sequence between themselves and a co-actor. I predicted that a decision-making actor will maximize the co-efficiency of the dyad by choosing an action plan that minimizes the overall costs of a sequence, given the available options. In a sequential object transfer task, participants made binary choices between paths to move along. The findings suggest a robust effect of total path minimization, providing initial evidence for co-efficiency.

Chapter 3 asks how people estimate the joint costs of an action sequence. I tested the hypothesis that when the costs of individual co-actors' actions are on the same scale, people compute the potential joint action costs as a weighted sum of the individual costs. Using an object matching task, I analyzed binary object choices as functions of individual costs and of their linear combination. Findings from three experiments show that participants minimize the combination of expected Self (decision-maker's) and Other (partner's) action costs.

Chapter 4 investigates if the co-efficiency hypothesis holds in contexts where different action types must be combined to achieve a shared goal. I predicted that people who consistently minimize a particular variable in individual action planning will also take into account a co-actor's

costs in joint action planning, even when individual costs are on different scales. Overall, I found support for this hypothesis, but participants focus more on minimizing their own costs than their partner's.

Finally, Chapter 5 tests the hypothesis that when interaction partners lack information about their partners' actions or an opportunity to communicate, co-efficiency might help coordination by being recognizable as a focal point. In an online version of the task, participants had to choose the same object as a remote partner, without feedback. I found a moderate effect of co-efficiency: some people recognize it as a potential focal point.

The findings of this dissertation together suggest that the expectation that social interactions unfold based on a principle of rational joint action is based in actual behavior. These studies open up new directions for research treating cooperative action planning in an economic framework.

Acknowledgements

This dissertation presents work undertaken in the past 5.5 years. I would like to express my gratitude to everyone who worked with me or supported me in any other manner during the years I spent at Central European University.

First of all, I am very grateful that my supervisors, Natalie Sebanz and Gergely Csibra agreed to take me and my project on, and that they supported me along the way with all matters small and large. I learned so much from them about how to conduct scientific research and how to present it to the world. Natalie and Gergő were always open to discussing new ideas and took the time to talk everything through so that we always came to a consensus about how to run things. I very much appreciate that we could work on this PhD together as a team of three, and I feel truly honored to have worked with them. Thank you for everything!

Secondly, this project includes studies that were done in collaboration with other young researchers from the department, from whom I learned plenty about research as well as about scientific collaboration. Study 1 was co-authored with Barbara Pomiechowska. Basia helped me a lot in terms of methods (programming) and also in writing my first paper for publication, which ensured that publishing as a student was not just a somewhat daunting, but also an exciting experience. I worked together with Oana Stanciu on Study 2. Oana took part in writing the paper, designing the experiments and analyzing the data we collected by lending us her expertise in statistical modelling. Needless to say, I learned a lot about this approach through conversation with her. To Basia and Oana: I am so happy to have worked with you, and I am grateful for the openness and curiosity with which you always engage in conversation.

In the SOMBY lab, I worked with many great colleagues. Crucially, I owe many thanks to Dávid Csűrös, Vanda Derzsi, Fruzsina Elekes, Fruzsina Kollányi, Fanni Takátsy, Eszter Salamon, Andrea Jenei and Anita Apor for their support in the lab and with all other issues related to running experiments. Without you, I would still be in the lab, setting up touchscreens... Thank you for all that you do to keep science going!

I would like to thank Christophe Heintz for challenging me on my project throughout the years. We may not agree on everything, but I greatly appreciate that you always took your time to engage with my work and to give me detailed feedback on it. Your point of view and the questions you raised informed the way I think about this project.

To Finta Györgyné Réka – a million thanks for all your hard work and endless support of us students.

I would like to thank the members of the Computational Social Cognition lab at Yale for warmly welcoming me to their group and for letting me participate in their inspiring and often hilarious lab life for a few months. I am especially grateful to Julián Jara-Ettinger for embarking with me on a project and guiding me whilst I was visiting in his lab.

The students and postdocs in the SOMBY lab and in the wider departmental community have been a pleasure to get to know throughout these years. So many of you I have grown to like a lot, and it would be impossible to list all of you, so please forgive me if I do not mention you by name. I will always think fondly of my PhD cohort: Francesca Bonalumi, Paula Fischer, Otávio Mattos, Helena Miton, and Simily Sabu, with whom we weathered many intense days and nights in the first year – and then the other years. ☺ I wish you all the best and thanks for the company!

Many hours were spent together in the Music Lab with Thomas Wolf, Denis Tatone, and Francesca Bonalumi, trying to find our hardest metal core, which injected much-needed glee and catharsis into the first part of this PhD. I cherish those memories. Likewise, I am happy we got to play ping pong on so many Friday evenings, it was the most fun way of releasing pressure. Thank you for the games, Gergő, Mikołaj, Barbu, and everyone else who dropped by!

And so are cherished countless memories of group outings, post-colloquia pogácsa-eating, dinners, conferences, winter and summer schools, corridor- and office-chats, coffee breaks, classes, workshops, protests, D&D sessions, nights spent clubbing, and other shenanigans. You were all integral parts of these experiences, in no particular order, and not only from CEU: Simily, Paula, Francesca, Laura Schlingloff, Zsuzsanna, Laura Schmitz, Merryn, Marcell, Denis, Basia, Thomas, Nazlı, Oana, Mia, Mathieu, Katka, James, Otávio, Nima, Barbu, Vjeran, Arianna, Nico, Luke, Kriszti, Sári, Márti, Berci, Rosie, Colin, Flora, Marlene, and Michael. I consider myself lucky for having spent time with you all.

Last but not least, I owe a special kind of gratitude to the family members and friends that kept my spirits up in the final one-year phase of this PhD – Bence Neszmélyi, Anna Török, Krisztina Török, Zsuzsa Metzger, Ágnes Tóth, Dóra Dulin, Merryn D. Constable, Márta Volosin, Berzsián, and Menyus. When I stopped believing in the chances of making it to the end, you encouraged me in conversations long and short (if you are a dog, that means you barked and were cute – well done!). Your support has been very important to me.

I would like to thank my parents for everything they have done for me.

Bence, thank you for always believing in me.

I dedicate this dissertation

to my sister, Anna.

Table of Contents

List of Abbreviations

2AFC	Two-Alternative Forced Choice
AIC	Akaike Information Criterion
AUC	Area Under the Curve
LOO-CV	Leave-one-out Cross-validation
PSE	Point of Subjective Equality
ROC	Receiver Operating Characteristics
WAIC	Watanabe-Akaike (or Widely Applicable) Information Criterion

Chapter 1. Introduction

The first part of *2001: A Space Odyssey* (Kubrick, 1968) depicts human evolution in striking images. It is heavily implied that a mysterious monolith inspires the ape-men to start using bones as weapons, triggering a process of evolution that eventually launches man into space. As compelling as this sequence is, it is quite unlikely that hominids evolved from bone-wielder to satellite-builder by the grace of an extraterrestrial object. Rather, humans have succeeded in establishing highly complex cultures thanks to their unique motivation and skillset for cooperation with their conspecifics (Tomasello et al., 2005; Tomasello, 2009). Engaging in cooperative activities and social learning has enabled humans to create and refine artifacts and behavioral practices over time (Tomasello et al., 1993), and to establish social institutions like marriage or money (Tomasello, 2009). A specific kind of cooperative activity is *joint action*, in which two or more actors work together to change the state of the environment in line with a goal that they share (Sebanz, Bekkering, et al., 2006).

Many everyday actions fall into the category of joint action: walking with a friend, playing table tennis or a concerto, writing an academic paper with colleagues, building a house. Such joint actions often differ from each other in terms of how much online coordination is involved between the movements of co-actors (e.g., in music-making [Wolf et al., 2020] versus in sequential object placing actions [Meyer et al., 2016]). They also differ in the ways a co-actor's intentions need to be represented (e.g., see minimal architectures [Butterfill, 2016] versus richer accounts of joint action [Bratman, 1992]). Accordingly, psychological research has approached joint action from multiple angles, exploring the mechanisms involved in online coordination and action planning (for reviews see Sebanz, Bekkering, et al., 2006; Sebanz & Knoblich, 2021).

Decision-making in the context of joint actions is a relatively recent topic in psychological research. For example, studies have investigated how interaction partners benefit from joint decisions (relative to individual decisions) in the integration of perceptual information when making visuospatial judgments about objects. Co-actors benefitted from having access to complementary information in the shared environment when they made spatial judgments on a continuous scale (Voinov et al., 2019). Furthermore, in binary decision tasks, the communication of metacognitive confidence information helped co-actors make optimal joint judgments about perceptual estimates (Bahrami et al., 2010, 2012), whereas factors like increasing group size hindered group performance (Sorkin et al., 2001). These studies probed, among other factors, the role of communication in collective decision-making, in contrast to research that tested non-verbal coordination using a behavioral economical task of convergence on an unknown target value (Roberts & Goldstone, 2011; see Sebanz & Knoblich, 2021). Smaller groups coordinated faster on the target, whereas larger groups adaptively differentiated the roles of their members to reach the goal (Roberts & Goldstone, 2011).

Collective, non-verbal decision-making about how to act next has been a less common research topic in the joint action literature. Some studies have explored what collective benefits may be derived from receiving information about the partner's actions and performance scores (Wahn et al., 2017) as well as the partner's gaze location (Brennan et al., 2008). When sharing these information, co-actors devised collaborative distributions of labor: they effectively divided the space between themselves in a visual search task (Brennan et al., 2008) and in a multiple object tracking task (Wahn et al., 2017), and achieved significant advantages in task performance. Additionally, continuous motor decision-making in interpersonal coordination has been investigated from a game-theoretical perspective, using tasks where co-actors' movements were directly coupled to each other (Braun et al., 2009, 2011). Lastly, Curioni and colleagues (2020) examined how people decide whether to engage in joint or individual actions to reach a given goal. They found that even at the expense of additional individual coordination costs (e.g., monitoring actions), people showed a preference for joint action (Curioni et al., 2020). My dissertation extends this latter line of research that probes decision-making in joint action planning without verbal communication, with a particular focus on the costs of actions.

In philosophy, deciding how to act is a part of practical reason/rationality (Bratman, 1987; Hampton, 1998; Wallace, 2020). The rationality of a person's decisions may be evaluated subjectively or objectively: the former examines if the goals chosen are coherent with the person's desires and beliefs, the latter whether the action taken was an optimal means of reaching the goal state (Stueber, 2006). Rationality from an objective point of view has been addressed most influentially by expected utility theory and game theory in economics and mathematics (Neumann & Morgenstern, 1944), and by statistical and Bayesian decision theory in cognitive science (Blackwell & Girschik, 1954; Körding & Wolpert, 2006; Trommershäuser et al., 2008). Practical rationality is crucial for navigating our social world as it provides a norm of interpretation of others' actions, summarized in the principle of rational action (Dennett, 1987; Gergely & Csibra, 2003).

The chapters that follow will present an experimental exploration of whether and how cooperating people make practically rational decisions that result in efficient instrumental joint actions, and if they do, whether this facilitates coordination under uncertainty. We tested the hypothesis that decision-making in a social setting ought to conform to the *principle of rational joint action*. This principle predicts that if people share goals in joint actions, they should also share and minimize the costs of actions that must be taken to reach said goals. We call decisions that instantiate the joint rationality principle and the resulting actions *co-efficient*.

This hypothesis is based on multiple strands of research. Accordingly, first I will describe why the rationality of actions is important for humans. On the one hand, it is central to making sense of others' actions (Csibra, 2017; Gergely et al., 1995), and on the other, it provides instrumental benefits for reaching goals (e.g., less effortful movements, Wolpert & Landy, 2012). Developmental and cognitive psychological findings suggest that humans expect others' actions to conform to the principle of rationality (section 1.1 The principle of rational action), and that these expectations are justified to some extent by actual behavior (section 1.2 Action planning as rational decision-making). Second, I will summarize relevant findings on joint action planning (section 1.3 Joint action planning) showing that co-actors represent their partner's actions and tasks in planning, and that they willingly incur additional effort themselves to reduce their partner's movement-related efforts. I will present the shared-effort model (Santamaria & Rosenbaum, 2011), which was proposed as an explanation of physical etiquette and served as a motivation for our studies. Finally, I will review relevant game-theoretic studies on interpersonal coordination (section 1.4 Game-theoretical accounts of interpersonal coordination).

We tested the co-efficiency hypothesis in four experimental studies (Chapters 2-5), the specific aims of which I summarize at the end of this chapter (section 1.5 Research Questions).

1.1 The principle of rational action

Understanding the social world around us is no easy feat. Since we do not have direct access to other people's mental states, their goals and intentions must be inferred based on the available sparse data. When we see a man carry a wooden box towards a tree, we might have trouble interpreting the observed behavior since multiple goals may be fulfilled by one action. He may want to open the box under the tree's shade to inspect its contents, to stand on top of the box to reach a branch, or to collect fallen leaves and store them in the box. How do humans solve the inductive problem of interpretation (Baker et al., 2009)?

They solve this problem by assuming that the man with the box is a rational agent. An agent's rationality is a condition for evaluating his or her goals and beliefs, which in turn enables the prediction of future actions (Dennett, 1987). The principle of rational action assumes that actions function to bring about desired goals, and that agents take the most efficient possible means to reach a goal under the constraints of the situation (Gergely & Csibra, 2003). This principle likely has a basis in the behavior of animals, whose actions tend to minimize energy expenditure and balance energetic costs against the benefits of foraging for food (for a review on how animals trade off different kinds of costs, see Cuthill & Houston, 1997). For example, horses' natural gait at any speed ensures that their oxygen consumption is minimized (Hoyt & Taylor, 1981); and some bird species' flight speed was found to maximize the efficiency of foraging[1] (i.e., the ratio of energy expended against the energy gained from food; McLaughlin & Montgomerie, 1990; Welham & Ydenberg, 1993). In sum, the principle of rational action states that agents maximize efficiency by minimizing the costs and maximizing the rewards related to an action.

1.1.1 Sensitivity to individual efficiency

Humans seem to have an early-emerging and pervasive expectation about the efficiency of observed actions. As infants, they readily attribute goals and intentions to agents (even when those lack human features, Gergely et al., 1995; cf. Heider & Simmel, 1944), which is crucial for social cognition – and goal attribution relies on assuming efficiency. The sensitivity to efficiency has mostly been studied in infants observing individual actions (Csibra et al., 2003; Király et al., 2003; Scott & Baillargeon, 2013; Skerry et al., 2013; Southgate et al., 2008). The infant is usually presented with animations of an abstract figure (e.g., Csibra et al., 2003) or live demonstrations of an adult agent (e.g., Scott & Baillargeon, 2013) performing means-end actions within various environmental constraints.

For example, an animated ball was depicted trying to catch another ball in Csibra and colleagues' study (2003). Looking times suggested that infants' expectation of efficiency had been violated when, following a change in environmental constraints, the actor did not adjust its movement path to be maximally efficient[2]. In contrast, they spent less time looking at events where the agent took the shortest available path to its goal. Infants can also infer an unachieved goal and environmental constraints based on observed actions that conform to the principle of rationality (Csibra et al., 2003), and they are able to predict an action based on knowing the actor's goal and the environmental constraints (Csibra & Gergely, 2007).

[1] Different optimal foraging theories emphasize the minimization of different costs. For example, the energy expenditure of an animal's young might also be part of a trade-off (Cuthill & Houston, 1997).
[2] Infants assume efficiency not just when it is operationalized as the shortest path towards a goal, but also as the shortest action sequence for obtaining a target (Scott & Baillargeon, 2013). Sixteen-month-olds expect that actors will choose objects that are reachable via an action sequence with fewer steps, and that are mentally more accessible – easier to see and keep in mind. Children from their second year of life are thus already sensitive to the efficiency of sequential actions, not just of single continuous movements.

1.1.2 Sensitivity to the efficiency of joint actions

Fewer are the studies that directly tested the sensitivity to the efficiency of observed joint actions, although there is evidence that infants understand that actors share goals in collaborative actions (Begus et al., 2020; Fawcett & Gredebäck, 2013, 2015; Henderson et al., 2013). Gredebäck and Melinder (2010) found that 6- and 12-month-olds could retrospectively evaluate the rationality of more or less rational social interactions, evidenced by pupil dilation measures. Babies expected feeding actions to be efficient and were surprised when they did not conform to this expectation.

Mascaro and Csibra (2014) addressed the computation of joint efficiency in 9- and 14-month-olds by systematically manipulating action costs. The infants were familiarized with animations of two abstract figures, where each of them passed tokens to the other through gaps in a simple maze's wall. These events were depicted as collaborations between the actors, and efficiency was operationalized as path length to the goal (i.e., the point of transfer between actors). During the familiarization phase, infants saw only jointly efficient actions, where the first actor always behaved individually inefficiently, but this inefficiency was justified by the environment's constraints and reduced the second actor's required effort – so that jointly, they acted efficiently. In test events, they either acted similarly jointly efficiently or jointly inefficiently. The jointly inefficient action still showed the first actor choosing a longer, individually inefficient path, but the environmental constraints changed so that they did not justify this individual inefficiency anymore. At 14 months, infants were sensitive to this joint inefficiency, and by implication, to joint efficiency (Mascaro & Csibra, 2014).

Building on this body of research, the rationality assumption was proposed as the main inferential principle guiding children's commonsense psychology (naïve utility calculus, Jara-Ettinger et al., 2016) and adults' core mentalizing abilities (Bayesian theory of mind, Baker et al., 2017). Jara-Ettinger and colleagues (2016) demonstrated in multiple experiments that toddlers make a variety of inferences based on observed actions: about an agent's preference for certain goals, a potential helper's competence and motivation to help, or past experience with an object.

For instance, when children observed two actors make the same low-cost choice, knowing the actors' subjective differences in rewards (preferences either for or against that low-cost choice) allowed them to infer the actors' competence for reaching different rewards (Jara-Ettinger et al., 2015). Younger infants can also infer the preferred goal of an agent: 10-month-olds expected that goals attained by costlier means must be more precious for the agent (Liu et al., 2017). These inferences are probabilistic; the hidden cost and reward functions of the agent are estimated through Bayesian inverse planning. In the case of adults, similar computational mechanisms are at play to enable the inference of beliefs and desires (Baker et al., 2017).

These theories of action understanding and social inference suggest that humans' social life relies a lot on the assumption of rationality. We hypothesize that the rationality principle also plays an important role in action planning in coordination contexts by making joint actions instrumentally efficient and co-actors more predictable to each other.

1.2 Action planning as rational decision-making

When people estimate probabilities in economic tasks, they make biased decisions[3] (Tversky & Kahneman, 1974; 1981). However, research on motor planning suggest that on the micro level, the motor system often works like an optimal planner: constantly trading off spatial and biomechanical costs in actions such as pointing, obstacle avoidance, reaching for and grasping objects (Elliott et al., 2004; Harris & Wolpert, 1998; Lyons et al., 2006; Todorov, 2004; Wolpert & Landy, 2012). Decisions about action conform to the "law of least effort" (Tolman, 1932) or "law of less work", stating that if two action sequences are equally reinforced, then the one requiring less work to attain the reinforcement will be learned (e.g., moving to food on a shorter path, Hull,

[3] Comparisons of choices in economic and motor decision-making tasks emphasized the importance of the different ways that probability information is provided for actors in each type of task (Maloney et al., 2007; Wu et al., 2009), or found that the two domains were not so different from each other (Jarvstad et al., 2013). Similar comparisons were also drawn between economic and sensory discrimination tasks (for a summary, see Summerfield and Tsetsos, 2020), and some argued that the different sources of uncertainty cause differences in behavior across domains (e.g., Juslin & Olsson, 1997). Summerfield and Tsetsos (2020) argue for an alternative view that emphasizes similarities rather than differences between domains, as well as the importance of efficient computations.

1943, p. 294). Such optimality is achieved by the prioritization of different types of costs in movement planning: costs intrinsic to the decision-maker – energetic (for a review, see Sparrow & Newell, 1998), biomechanical (e.g., Rosenbaum et al., 1990), or cognitive (e.g., Fournier, Coder et al., 2019; Kool et al., 2010) –, and extrinsic costs such as money (Trommershäuser et al., 2008).

1.2.1 Intrinsic action costs

On the one hand, the planning of pointing actions has been found to minimize energetic costs (Lyons et al., 2006; Elliott et al., 2004). Lyons and colleagues (2006) pointed out that primary movement segments are usually not centered to a target, and actors have an undershooting bias when their hands stop short of the target location. The magnitude of this undershooting bias depended on the degree to which movement direction was congruent or incongruent with gravity. Moving downward, actors exhibited larger undershooting than moving upward, because the relative energetic costs of overshooting and making corrective movements against gravity would have been larger than doing it *with* gravity (Lyons et al., 2006). The modulation of this bias suggests people minimize energetic costs in pointing. On the other hand, spatial error may also be minimized in pointing (Battaglia & Schrater, 2007; Harris & Wolpert, 1998).

The type of costs that are traded off in planning may depend on the movement context. For example, the relative weighting of spatial error and biomechanical costs in manual obstacle avoidance was found to depend on sensory uncertainty, motor noise, and practice (Cohen et al., 2010). When people had to clear an obstacle with their hands while holding a dowel, high visual uncertainty and motor noise favored the prioritization of spatial error minimization (movement clearance was larger than in low-uncertainty and low-noise trials). Furthermore, with practice, low uncertainty and low noise, participants minimized their biomechanical costs by reducing clearance so that their limbs moved in a smaller range (Cohen et al., 2010).

Planning trades off naturalistic costs of goal-directed action when selecting actions on a higher functional level, such as choosing hand grasps. The *end-state comfort effect* suggests that actors aim to minimize biomechanical costs in action selection (Rosenbaum et al., 1990; Cohen &

Rosenbaum, 2004). When participants moved a vertical cylinder from a home platform to a target platform of varying heights, an inverse relationship was found between grasp height and target platform height. The lower the cylinder had to be placed, the higher actors grasped it (Cohen & Rosenbaum, 2004). These findings suggest that when grasping an object, the final position of the hand is adapted to accommodate future task demands, which often entails that the initial position of joints will be awkward (Rosenbaum et al., 1990, Rosenbaum et al., 1996).

Actions usually also entail cognitive costs, for instance, related to memory retention. The *precrastination* effect, first shown by Rosenbaum and colleagues (2014), describes people's tendency to expend additional physical effort to hasten the completion of a subgoal in an action sequence. In multiple experiments, people preferred to pick up a loaded bucket early along a walking trajectory, rather than later, even though the latter solution would have ensured a shorter load-carrying distance. The authors suggested that such behavior is rational if we account for the cognitive costs of keeping a subgoal ("pick up the weight") in working memory, beside the costs of movement effort (Rosenbaum et al., 2014). Subsequent studies proposed that people precrastinate to start, rather than to finish a sub-task sooner, to save cognitive resources (Fournier, Coder, et al., 2019; Fournier, Stubblefield, et al., 2019), and that precrastination is less likely to happen when it would incur further cognitive costs (Raghunath et al., 2020). Precrastination reflects a trade-off between cognitive and physical action costs.

1.2.2 Extrinsic action costs

Beside intrinsic costs, motor planning trades off extrinsic costs. Actors were found to take into account monetary costs and rewards during planning, and maximized gains for rapid, ballistic pointing movements (Trommershäuser et al., 2003a, 2003b). Trommershäuser and colleagues (2003a) asked participants to point, under time constraint, at a touchscreen where reward and penalty regions were labelled. The location of aiming shifted depending on the changes of the environment's cost-reward structure and mean aiming points were close to the optimal point that maximized the expected gain of movement. Human performance compared well with an optimal planner model (Trommershäuser et al., 2003a, 2003b).

To summarize, the literature on action planning as decision-making provides evidence that the principle of rational action is usually justified in the case of individual actions: people tend to plan and execute goal-directed actions in an efficient manner, by flexibly trading off different kinds of costs against one another. Is this the case for planning joint actions, too?

1.3 Joint action planning

Joint action research has so far not provided studies either to directly address rational planning in coordination in general[4], or specifically to test if people minimize the joint costs of two co-actors' coordinated action sequences when they choose between multiple joint action plans. However, behavioral and neurophysiological findings suggest that people represent their partner's anticipated actions and the constraints of an interaction partner's task (providing cues to the action costs) in joint action planning. These findings come from experiments that required participants to simultaneously or sequentially coordinate their actions, and in the latter case the representations of a partners' actions and task constraints were often manifested in helping behaviors.

1.3.1 Planning with the co-actor in mind

Comparisons of event-related potentials (ERP) in the preparatory phase preceding action in individual and joint tasks showed that actors represent their partners 'actions in a functionally equivalent way to their own (Hommel et al., 2001). This was reflected in increased inhibition of a tendency to act in turn-taking tasks with low coordination demands (Sebanz, Knoblich, et al., 2006; Tsai et al., 2008, 2006), because the partner's goals activated the motor representations of actions in the actor, too, who had to wait for his/her turn in the task. Furthermore, when a task required precise temporal coordination between co-actors, an increased amplitude of the contingent

[4] Note that since the start of our project, Curioni and colleagues (2020) ran a study where they looked at decisions to coordinate (or not) with others as a function of potential individual utilities of actions in individual and joint contexts.

negative variation (CNV), a component reflecting action planning, suggested that participants prepared both their own and their partner's part of the task (Kourtis et al., 2014). Beyond keeping track of the specific tasks of each co-actor and simulating a partner's future actions (Kourtis et al., 2013), people form we-representations that specify the outcome of a joint action on the group level, not just on the level of individual contributions (Kourtis et al., 2019; see also Della Gatta et al., 2018; Gallotti & Frith, 2013; Sacheli et al., 2018). This was apparent in the modulation of EEG indices (P600, alpha suppression, CNV) related to cognitive and sensorimotor representations of the relationships between interaction partners' upcoming actions in a pre-cueing task (Kourtis et al., 2019), and in collective goal-related interferences that were observed in line and circle drawing tasks (Della Gatta et al., 2018).

Behavioral studies showed that, in addition to representing upcoming actions, people represent others' task constraints. For example, the height of an obstacle to be cleared by a co-actor's hand (Schmitz et al., 2017) and the order of the sub-goals in a co-actor's designated action sequence (Schmitz et al., 2018) influenced the kinematics of an actor's hand movements in simultaneous joint tasks. Participants deviated from their individually most efficient movement paths, which helped synchronization between co-actors (Schmitz et al., 2017). Conversely, in other studies' sequential joint tasks, although it was not necessary for the participants to synchronize their movements, the facilitation of a partner's actions suggested that the partner's task constraints were co-represented and integrated into action planning.

In sequential joint object manipulation tasks, people adjusted their own actions to reduce the effort of their partner who concluded the action sequence, either by rotating objects for them (Constable et al., 2016; Dötsch & Schubö, 2015; Ray & Welsh, 2011), selecting a hand and grasp type (Scharoun et al., 2016), choosing grasp locations on an object (Meyer et al., 2013), or choosing appropriate spatial locations for handing over objects to the partner (Gonzalez et al., 2011; Scharoun, et al., 2017; Ray et al., 2017). Dötsch and Schubö (2015) found that in a joint pick-and-place task requiring the rotation of an object, the first actor systematically altered her grasp and rotation angles to reduce her partner's predicted effort in completing the action sequence. Participants incurred these additional individual costs without being instructed to do so.

1.3.2 The shared-effort model of coordination

As several authors noted (Meyer et al., 2013; Dötsch & Schubö, 2015; Ray & Welsh, 2011; Scharoun et al., 2016, 2017), a possible explanation for these facilitatory behaviors is Santamaria and Rosenbaum's (2011) *shared-effort model.* The model posits that in a social context, when people have a choice either to coordinate with others to perform an action or to complete it independently, they will coordinate to reduce shared effort (Santamaria & Rosenbaum, 2011).

In an observational study of door-holding behavior, Santamaria and Rosenbaum (2011) found that the probability of holding a door open for others was proportional to the followers' distance from the door, and the more people followed, the longer a person held the door open for them. Followers sped up to reach it, suggesting that they also tried to reduce group effort. Based on these findings, the authors suggested that if a person at the door 1) thinks that by holding it open for others, effort costs can be relatively minimized – i.e., aggregate group costs are lower than the sum of individuals' costs if acting alone – and 2) believes that the follower shares her belief, she will opt to hold the door open even at the expense of exerting additional individual effort (Santamaria & Rosenbaum, 2011). Framing this in terms of rationality, coordination is based on a decision to maximize aggregate expected utility by sharing the costs[5]. The prediction follows that, if an actor finds that the potential shared action cost in coordination would be lower than the sum of each individual's costs, she will expend additional efforts to coordinate.

However, often we need not only decide whether or not to coordinate, but to decide on which course of action to follow together. Previous research on the representation and accommodation of a co-actor's effort in joint action does not answer the question whether a

[5] Although Curioni and colleagues (2020) did not consider this as a possible explanation for their results and they did not calculate joint utilities, theoretically, Santamaria and Rosenbaum's model (2011) could be one potential account for their results. They demonstrated people's propensity to decide to coordinate with another person, even at additional individual costs.

"principle of rational joint action" guides action planning in coordination tasks where actors already share a goal and try to decide how to act. According to this principle, cooperation would entail the minimization of shared action costs, rather than the minimization of only the individual costs of the decision-maker or the partner (i.e., facilitation). Therefore, to test the co-efficiency hypothesis, it is necessary to examine decision problems in which both kinds of behavior would contribute to co-efficiency, as a function of the environmental constraints. However, the aforementioned paradigms only addressed situations where, from a joint perspective, the individual *should* have incurred additional individual costs to reduce joint costs (as in holding a door open, Santamaria & Rosenbaum, 2011). Theoretically, in other cases, the joint action planner should refrain from taking over some of her partner's costs to minimize shared effort and therefore facilitation of the partner's efforts would be detrimental to co-efficiency. This gap in the literature was our primary motivation for testing the co-efficiency hypothesis.

1.4 Game-theoretical accounts of interpersonal coordination

Minimizing effort in motor interactions has also been investigated in a game-theoretic framework (Braun et al., 2009, 2011; Grau-Moya et al., 2013). These studies focused on the strategic aspect of movement control, using continuous motor versions of the prisoner's dilemma, a rope-pulling game (Braun et al., 2009), and classical coordination games (e.g., stag hunt, Braun et al., 2011). Payoffs were defined in terms of effort due to resistive forces against hand movement. People defected more in the motor prisoner's dilemma than in classic economic versions of it (Braun et al., 2009), but generally converged on cooperative solutions in the coordination games (Braun et al., 2011). These studies employed simultaneous coordination tasks in which the participants' movements were directly coupled either with another human's (Braun et al., 2009, 2011) or a virtual partner's (Grau-Moya et al., 2013). That is, the actions of one actor directly affected not just his or her own payoffs, but the co-actor's payoffs, too, and vice versa. The authors argued that the observed patterns of coordination or competitive Nash equilibria originated from the dynamical coupling of the sensorimotor processes of interaction partners (Braun et al., 2009).

Game-theoretic studies such as these provide a promising new methodology for research on cooperation and joint action by treating interactions within an economic framework (see also Engemann et al. [2012] for an argument for using stag hunt games in neuroscientific research on cooperation). Our project complements these works in two ways: we aimed to explore how the cost-reward landscape of the environment influences co-actors' discrete choices between joint action plans, without information about a partner's motor processes through direct coupling.

1.5 Research Questions

This dissertation aims to contribute to work addressing motor planning as rational decision-making on the one hand, and to joint action research investigating how cooperative actors represent their co-actors' task constraints when they plan in different coordination contexts on the other hand. Our goal is to provide an account of rational joint action planning that connects the topics of practical rationality and interpersonal coordination by adopting an economic approach to planning.

Study 1 (Chapter 2), as a proof of concept, tested the prediction of Santamaria and Rosenbaum's (2011) shared-effort model pertaining to the minimization of aggregate effort (we did not probe the role of sharing beliefs about costs in the occurrence of the behavior). We investigated the question whether, adults maximize the co-efficiency of the dyad when coordinating in a sequential joint task. At the same time, we aimed to extend Santamaria and Rosenbaum's (2011) hypothesis to situations in which people already share a goal and coordination is a given. Based on their observations, we predicted that co-actors who share a goal in a sequential joint task will aim to maximize co-efficiency, even by incurring additional individual costs, if necessary. We operationalized maximizing efficiency as taking the shortest available path to a goal given the environmental constraints, consistent with research on the teleological stance in infants (Csibra et al., 2003), which provided ample evidence that observers are sensitive to such visual proxies of action costs. Participants transferred a football on a touchscreen between each other and made binary choices between more or less co-efficient action plans. We analyzed their decisions in four experiments[6].

How do people compute the joint costs of action sequences? We addressed this question in *Study 2 (Chapter 3)*, which adopted a parametric experimental approach. The integration of individual action costs in action planning appears to us as a conceptually similar problem to perceptual cue integration. People optimally combine cues from haptics and vision (Ernst & Banks, 2002), and – provided that judgments are weighted according to their reliability –, they can also combine perceptual judgments across interaction partners to achieve optimal collective decisions (Bahrami et al., 2010). In the action domain, a computational model of competition and cooperation operationalized cooperative planning as a maximization of a joint utility function (Kleiman-Weiner et al., 2016). Based on this, we hypothesized that the computation of joint costs will rely on a weighted summation of individual action costs when those are easy to estimate (i.e., they are on the same scale: distance). We employed a new coordination task, an object matching game, that allowed for the independent manipulation of individual action costs in a stochastic fashion, and in which joint-cost minimization was only possible by summing the individual action costs. Using a hierarchical Bayesian regression model, we analyzed the decisions of participants and compared them to the predictions of strategies that minimized either the decision-maker's own, the partner's, or the dyad's action costs. We also tested an alternative strategy that prioritizes fairness in between-partner cost distribution over efficiency.

Study 3 (Chapter 4) tested whether co-efficiency generalizes to more naturalistic joint action contexts. In real-life social interactions, people cooperate in activities that are composites of

[6] This study has been published and Study 2 is currently under review for publication. The published and submitted Supplementary Materials accompanying these studies are presented right after the main texts of these papers, to aid understanding of the multiple experiments reported in these chapters. In contrast, longer Appendices for Chapters 4-5 follow after the Discussion, titled Appendix A and B, respectively.

different but complementary individual actions. For example, meeting up by the park bench for a chat may inspire me to take public transport before walking in the park to the agreed location, whereas you might decide to ride your bike from your home to the park before walking to the bench to greet me. Most of the time, there are multiple ways to solve an individual or a joint task by combining different actions in sequences. Building on research that quantified judged relative action costs in individual action selection tasks (e.g., Potts, Callahan-Flintoft, et al., 2018; Rosenbaum et al, 2011), we used psychophysical curve-fitting to estimate individuals' preferences for each type of action employed in Study 3's task. The task was a modification of the object matching task in Study 2, and participants again made binary decisions to act. Using the judged relative costs of tapping and dragging actions on a touchscreen, we estimated the joint action costs of composite sequences. Assuming that people maximize efficiency in an individual decision task, we tested the hypothesis that they would integrate their partners' action costs in joint action planning as well – even when those costs are on a different scale from their own.

Finally, *Study 4 (Chapter 5)* addressed a potential function of the principle of rational joint action unrelated to the efficiency of face-to-face coordination: it might make interaction partners more predictable to each other when trying to coordinate under uncertainty about the other's actions. In theory, expecting that an interaction partner will make a mutually beneficial decision can help in contexts where selfish or altruistic decisions may result in a failure to coordinate. This study drew from the literature on coordination games in game theory and behavioral economics, specifically from the theory of focal points that provide solutions to coordination problems due to their saliency (Schelling, 1960). It is unclear whether payoff combinations with asymmetries between the individual payoffs (which often characterizes co-efficient actions) are used by people as focal points, and we tested whether co-efficiency helps people in a coordination game. We predicted that people might find the co-efficient action plan salient as it better restricts the range of potential focal points than other features in a task, and it is a utility-maximizing option. In an online object matching task adapted from Study 2's task, we asked participants to guess the trial-by-trial decisions of a remote partner. Object choices were analyzed in terms of potential salient features such as shape or relative location on the screen (left vs right side etc.), and the frequency of successful coordination was compared against chance. We analyzed decisions to see if the minimization of the co-actors' joint costs was recognized by participants as a potential coordination strategy.

Chapter 6 provides a discussion of the findings of these four studies, summarized in light of the literature described above. I will also discuss open questions that our studies raise for future research, for example, regarding the boundary conditions of co-efficient decision-making. Finally, I draw conclusions on what our findings mean for the extension of the principle of rational action to social interactions.

Chapter 2. Rationality in joint action: Maximizing co-efficiency in coordination

2.1 Introduction

People tend to act efficiently when they aim to achieve a goal. For example, on a shopping visit to a mall, shoppers keep to the minimum the walking distance covered between shops of interest (Gärling & Gärling, 1988), trying to get what they need with the least effort. Motor planning of everyday gestures and movements, such as pointing and grasping, follows the same principle. People move with minimum effort when pointing (Lyons et al., 2006), and guide the movement of their hand to ensure a stable grasp at first contact and to minimize post-contact adjustments (Christopoulos & Schrater, 2009). Furthermore, people sometimes adopt uncomfortable hand positions when these are helpful to continue their action after retrieving an object, suggesting that they plan actions with the total expected effort in mind (Cohen & Rosenbaum, 2004). The motor system often performs comparably to an optimal decision-maker (Wolpert & Landy, 2012), selecting the most beneficial solutions in the given circumstances.

How do people achieve efficiency when they work together? Joint actions are aimed towards accomplishing shared goals and require coordination between two or more partners (Butterfill, 2016; Sebanz, Bekkering, et al., 2006). If each interaction partner were to maximize the efficiency of their individual actions, this could lead to sub-optimal joint performance or a failure to coordinate. Imagine that two friends spot each other from the two ends of a park and would like to sit down for a chat. If each of them walked to the bench closest to her, minimizing her individual cost in terms of walking distance, they may end up sitting on different benches. Sharing the benefits of achieving a joint goal may demand from the actors to share the costs as well. Importantly, there are multiple ways to do so, depending on whose costs they want to minimize. How do people distribute the costs of joint actions?

Accounts of team reasoning have proposed that people maximize the aggregate benefits and minimize aggregate costs of the group (Gilbert, 1987; Hurley, 2005; Sugden, 2000), and empirical evidence for these claims has been provided through interactive economic games (e.g., Colman et al., 2008a). Minimizing aggregate, rather than individual, costs of an action for a fixed benefit entails aiming for 'co-efficiency', rather than individual efficiency.

Recent studies have shown that people facilitate their partner's performance by reducing the partner's costs. In tasks where participants handed over objects to another person, they adjusted their own actions to reduce the effort of the partner who concluded the action sequence. They rotated objects (Constable et al., 2016; Dötsch & Schubö, 2015; Ray & Welsh, 2011), selected particular grasp types (Scharoun et al., 2016), chose appropriate grasp locations on an object (Meyer et al., 2013), and handed over objects at spatial locations that made it easier for the partner to finish the task (Gonzalez et al., 2011; Ray et al., 2017; Scharoun, et al., 2017).

Further evidence for spontaneous sharing of effort comes from an observational study that investigated how people hold doors open for others behind them (Santamaria & Rosenbaum, 2011). The closer a follower, the more likely people were to hold open the door; the door was held open for longer when two people followed than when only one followed; and when the door was held open, followers sped up to reach it. While these findings are generally in line with the idea that people are sensitive to aggregate group effort, they do not clarify why. People might be helping their partners; that is, people might incur extra costs to reduce the partner's costs. Alternatively, people might act co-efficiently, which differs from altruistic behavior in that the person incurring costs aims to minimize aggregate group costs rather than the co-actor's costs.

Numerous real-world situations, from cooking together through dividing paperwork to raising children, require partners to coordinate and invest efforts to achieve shared goals. To shed light on the question of how people distribute costs of joint activities, we pit co-efficiency against helping by investigating co-efficiency of joint action planning in the context where individual and aggregate costs of two actors were systematically manipulated. We operationalized action cost as proportional to path length in a task that required participants to move objects from one location to another. In this context, maximizing efficiency amounts to taking the shortest available path to a goal, given environmental constraints. The joint version of the task involved passing an object to a partner at one of two transfer locations (Figure 2.1). The person passing the object could optimize either her own efficiency, choosing the shortest sub-path to a transfer location, or the total executed path length of the dyad. In some trials, taking the shorter sub-path from an individual perspective resulted in an overall shorter path for the dyad (congruent trials). In other trials, taking the shorter sub-path from an individual perspective corresponded to an overall longer path for the dyad (incongruent trials). In further trials, the two paths were equal in length from a dyadic point of view (neutral trials) but differed in terms of the relative sub-path lengths of the two actors. If people maximize co-efficiency, they should specifically incur higher individual costs on incongruent trials to reduce joint costs. If they maximize individual efficiency, they should consistently take the shorter sub-path, regardless of the overall joint costs. Finally, if people are being helpful, they should act to minimize their partner's individual cost, either only when this does not impair co-efficiency (on neutral trials) or even when it does (taking the longer sub-path on congruent trials, which would minimize the sub-path for the partner, but increase the overall path length).

To ensure that the costs associated with the different paths are perceivable and that our task affords cost optimization, we first ran an individual version, where single participants performed both steps of the object transfer task (Experiment 1). We then investigated joint performance (Experiment 2).

2.2 Experiment 1: Individual Efficiency - Methods

This experiment tested whether people maximize efficiency of individually executed action sequences. We gave participants a choice between two paths along which they could move a ball: shorter vs. longer. If people act efficiently, they should consistently select the shorter path. The exact proportion of efficient choices might be influenced by the degree of asymmetry between available paths: the larger the length difference between the paths, the more sensitive people might be to cost differences. To test this, we manipulated the difference in length between paths.

2.2.1 Participants

Target sample size was determined by power analysis in G*Power 3, for a medium effect size (d = 0.6) on binary choices with a one-sample t-test against a 50% chance level, using an α = .05 (Faul et al., 2007). A sample size of N = 24 was estimated to provide 80% statistical power. The participants were recruited through Central European University's Research Participation System (SONA Systems) and a student job agency. They gave their informed consent and received vouchers in exchange for their participation. The study was approved by the United Ethical Review Committee for Research in Psychology (EPKEB) in Hungary. Twenty-seven right-handed participants took part in Experiment 1. We analyzed the data of 24 participants (7 males, age M = 25.1 years, SD = 3.54). We excluded three additional participants due to a computing error (n = 1) or an experimenter error (n = 2).

2.2.2 Apparatus

The task was performed on a horizontally placed touchscreen monitor (Elo Touch, 2201L, 22", resolution 1920 X 1080 pixels, 60 Hz) connected to an Apple iMac computer. Stimulus presentation and data recording were controlled by a script using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997; Kleiner et al., 2007) in MATLAB® (The MathWorks, Natick, MA). A response box (Black Box Toolkit Ltd.) was used to control trial onset.

2.2.3 Stimuli

On each trial, the display consisted of the following elements: the image of a football, a starting position, a goal position, and obstacles (Figure 2.1). The starting and goal positions, marked by squares, were located in diagonally opposite corners of the screen. The ball was initially placed at the designated starting position. Obstacles consisted of (1) a wall placed in the middle of the screen, separating the starting and the goal position, with two gaps to pass the ball (marked by circles), and (2) two additional barriers, located perpendicularly to the wall on each of its sides

(Figure 2.1). One barrier had a fixed (maximal) length of 1 unit and was located either on the side of the screen nearer to the participant, or on the side farther from her. The size of the barrier on the other side of the wall varied between 0 (no barrier) and the maximal length in 0.25 unit steps, resulting in five distributions of barrier lengths: 1–0, 1–0.25, 1–0.5, 1–0.75, 1–1. These combinations of barrier lengths on each side of the screen provided the participants with different degrees of asymmetry between the costs of moving to the gap closer to or farther from their starting positions. For example, a barrier of 0.75 unit length on the participant's side resulted in a much longer sub-path to the gap farther away from them than the sub-path to the gap closer to them. In contrast, a 0 unit long barrier (i.e., no barrier perpendicular to the wall between the two sides of the screen) imposed the least difference between the short and long sub-path options for the participant. The orientation of the wall with the circled gaps in it was either parallel or perpendicular to the longer side of the touchscreen, with half of the trials displaying a horizontally, the other half a vertically oriented wall.

2.2.4 Procedure

The starting position of the football was always on the side of the participants. They were instructed to pull the ball with their finger from the starting position to the goal position through one of the gaps in the wall. The movement of the ball was blocked if any pixel of the ball image overlapped with a pixel of the displayed walls, barriers and screen boundaries – an event we will refer to as a collision. All instances of such collisions were registered and signaled to the participants by an audio soundbite. Participants were instructed to complete the task as accurately as possible, i.e., with the least amount of collisions with the obstacles.

The participants were instructed to keep their dominant hand on the response box at the beginning of each trial. The box was placed perpendicularly along the middle of the longer side of the touchscreen. This ensured that the key on the box was equidistant from the two potential starting positions at the left and right corners of the screen. Once the participants started pressing the key on the response box, the layout was presented without the ball. After 1500 ms, the ball appeared in the starting position, which prompted the participants to release the key and start moving the ball. When the ball arrived in the circle at one of the gaps, the sub-goal was completed. To indicate this to the participants, the background of the circle was highlighted, the movement of the ball was blocked, and the participants had to briefly release it before they could resume dragging it further. As soon as the ball arrived at the goal area, a short auditory signal marked the completion of the trial.

Before the experiment, the participants completed a brief practice session of 10 trials to familiarize themselves with the use of the touchscreen, the response box, and the screen layouts. They then completed 80 trials: 32 congruent, 32 incongruent, and 16 neutral trials. In congruent trials, passing the ball through the gap closer to the starting position (i.e., taking the short sub-path to the sub-goal of passing through the wall) coincided with taking the overall shorter path to the goal location (Figure 2.1a). In the incongruent trials, the short sub-path was part of the longer total path to the goal location (Figure 2.1b). Neutral trials were symmetric in terms of total path lengths (Figure 2.1c). The length of the shorter barrier in the non-neutral trials, the orientation of the layout (horizontal or vertical wall), and the starting positions (left or right side) were fully counterbalanced. The order of trials was randomly determined. Participants completed the task on average in 14.22 minutes (*SD* = 2.11). At the end, participants filled in a short questionnaire on what they thought to be the purpose of the experiment.

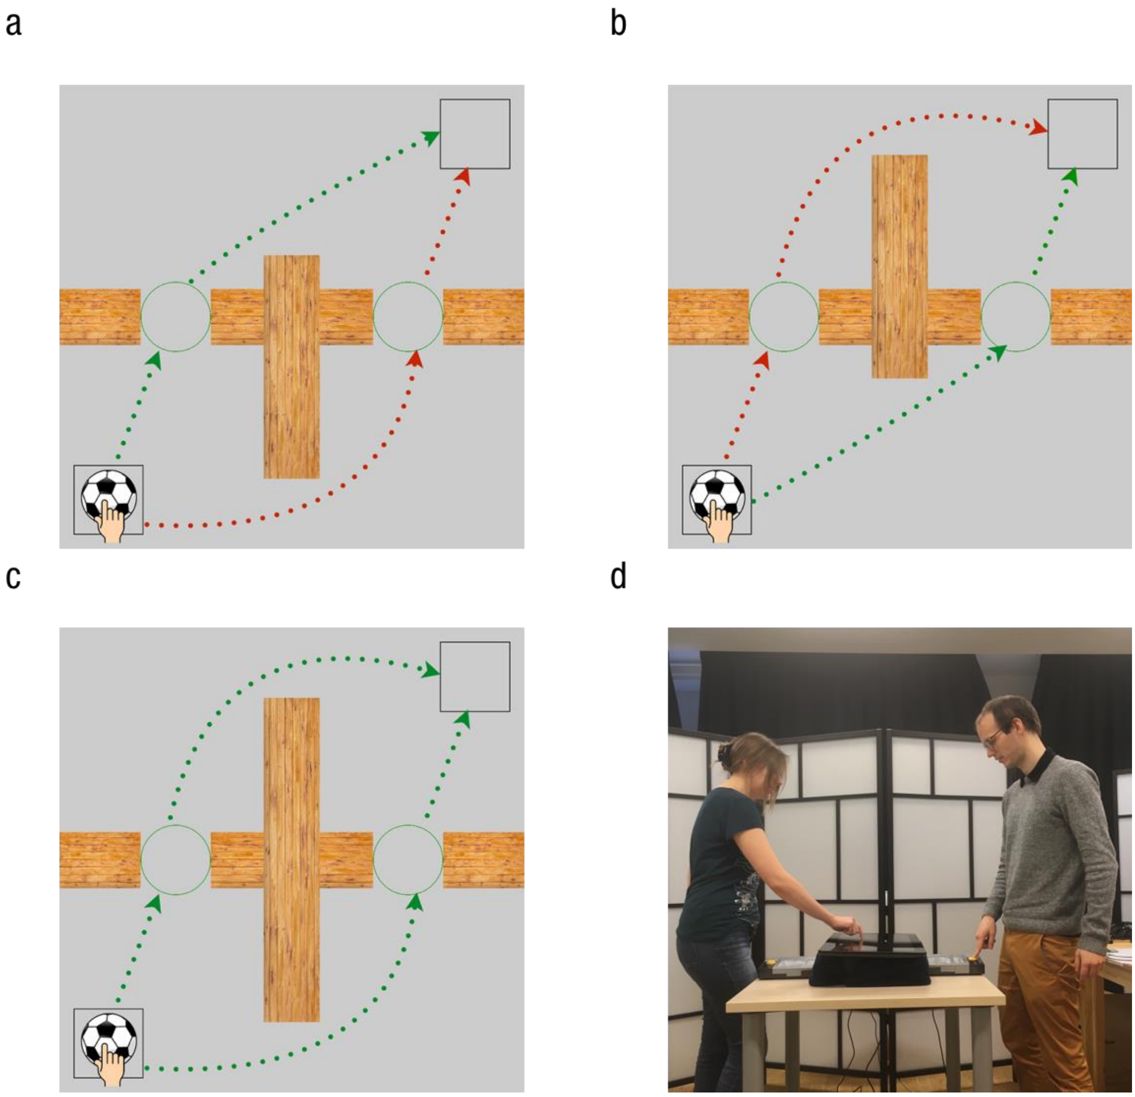

Figure 2.1: Examples of layouts used in the task. The participants' task was to move the image of the football between the squares that indicated starting and goal locations, through one of the circles that marked the possible sub-goals. Layouts belonged either to the (a) Congruent, (b) Incongruent or (c) Neutral condition, depending on whether choosing the shorter sub-path to a sub-goal resulted in a (a) shorter, (b) longer, or (c) equal total path length compared to the other option. Efficient (Experiment 1) and co-efficient (Experiment 2) total paths are colored green, sub-efficient paths are colored red (the arrows in the figure are just for illustration). (d) The experimental setup and the actors' positions in Experiment 2.

2.2.5 Data analysis

The primary dependent variable was the proportion of efficient path choices, that is, the proportion of trials where the participants chose the shorter total path between the starting and goal locations. Choice proportion data were not normally distributed, therefore all statistical analyses were performed on arcsine transformed proportion data. All comparisons were conducted in JASP (JASP Team, 2018) using Wilcoxon signed-rank test (two-tailed), unless otherwise noted. We report V statistics for the Wilcoxon tests, as well as matched-pairs rank-biserial correlation coefficients (*r*), both provided by JASP. The *V* statistic corresponds to the sum of ranks assigned to positive-signed differences between the two tested paired samples and represents the value to be compared to those found in tables for Wilcoxon test. The matched-pairs rank-biserial correlation coefficient (*r*) represents the effect size of the difference between the paired variables. The lower the value of *r*, the lower the difference between positive and negative rank sums, therefore the smaller the size of the effect that rendered the two paired samples different.

To assess whether choosing the efficient option resulted in faster or more accurate performance, we also analyzed mean number of collisions per trial (to estimate accuracy) and total trial durations (to estimate average speed) according to the choices actors made. Duration measurements were log-transformed for analyses. For ease of reading, text and figures report untransformed summary statistics.

2.3 Results

2.3.1 Proportion of efficient choices

Participants tended to minimize the total path length. They transferred the object in an efficient manner, i.e., through the gap that was closer to them in the Congruent trials (M = .88, SD = 0.21), and through the farther gap in the Incongruent trials (M = .80, SD = 0.28) (Figure 2.2a). Efficient choice ratios were significantly different from chance (Congruent: V = 294, p < .001, r = .96, 95% confidence interval (CI) for the efficient choice proportion difference from chance level [arcsine transformed chance level of 0.5 = 0.7854] = [1.21, 1.48]; Incongruent: V = 253, p < .001, r = .69, 95% CI = [1.06, 1.42]). Efficiency did not differ between the Congruent and Incongruent trials, as suggested by a paired-samples comparison between the ratio of efficient choices in the two conditions (V = 116, p = .065, r = -.23, 95% CI = [-0.02, 0.26], see Figure 2.2a). In the Neutral trials, participants tended to choose the closer gap (M = .67, SD = 0.18; V = 234.5, p < .001, r = .56, 95% CI = [0.89, 1.09]). Paired-samples comparisons to matching sub-path choices in the Neutral condition showed a significant increase in the proportion of closer-gap choices in

Congruent trials (V = 239, p < .001, r = .59, 95% CI = [0.26, 0.52]) and a significant decrease in Incongruent trials (V = 4, p < .001, r = -.97, 95% CI = [-0.80, -0.45]). That is, compared to the Neutral trials, in asymmetric trials participants shifted their decision towards the more efficient choice.

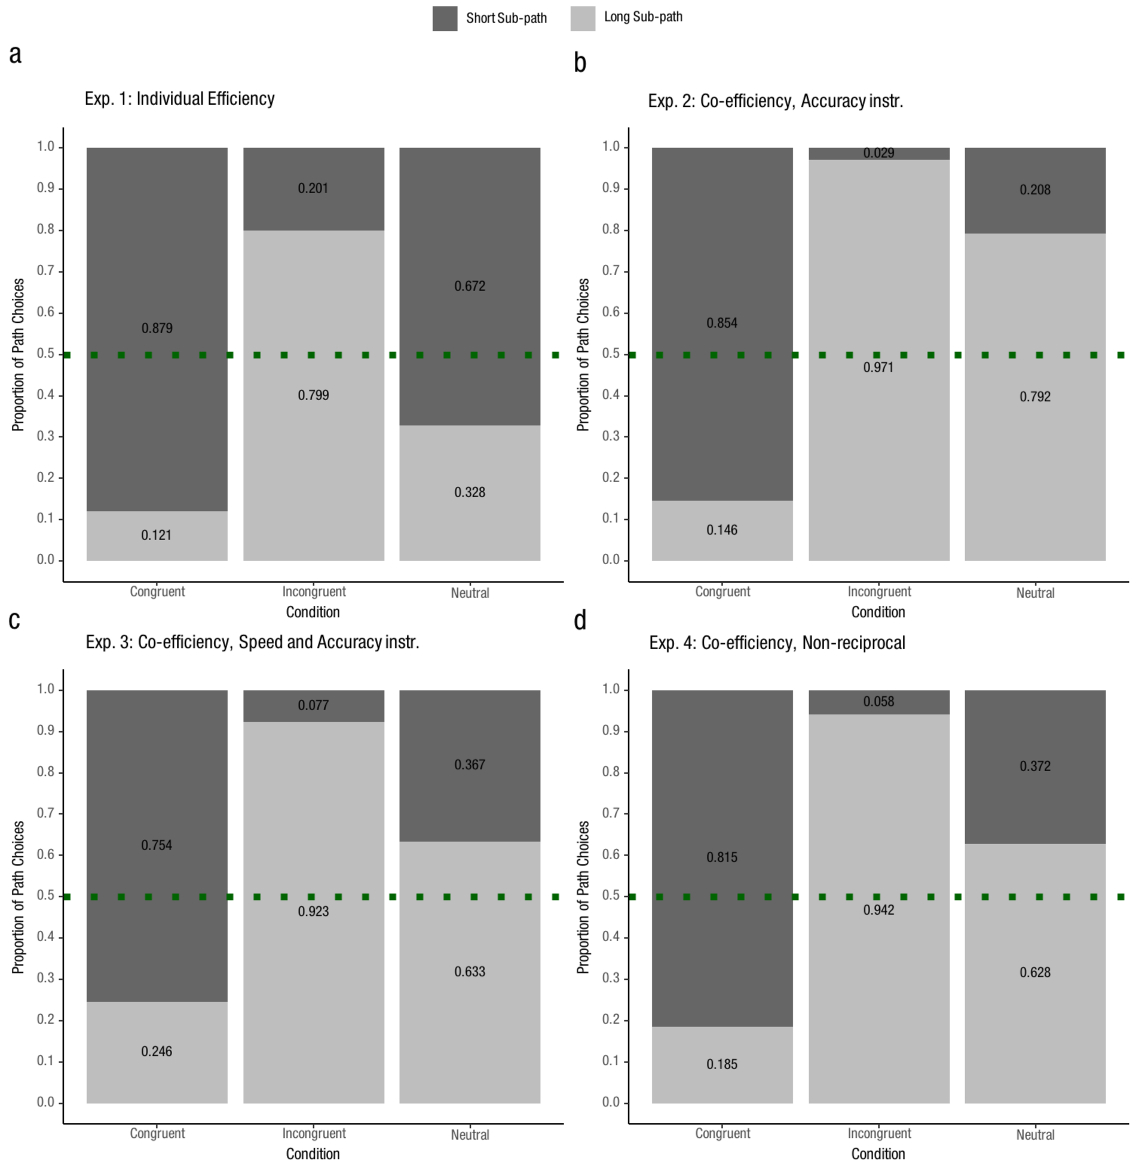

Figure 2.2: Illustrations of means of raw path choice proportions in the three conditions (N = 24) of Experiments (a) 1, (b) 2, and two additional joint experiments: (c) Experiment 3, and (d) Experiment 4. Experiments 2 and 3 differed only with regards to the task instructions, whereas in Experiment 4, only one of two partners made choices. Further details and results of Experiments 3 and 4 may be found in the SOM-R. Efficient (Experiment 1) and co-efficient (Experiment 2-4) choices were taking the short sub-path (dark grey) in the Congruent condition, and the long sub-path (light grey) in the Incongruent condition. Dotted lines indicate chance level (0.5) of choice proportion.

We analyzed whether the size of the difference in length between the path options had an effect on participants' efficient path choices using a 4 X 2 repeated measures ANOVA with Cost

Asymmetry (0, 0.25, 0.5, 0.75 unit lengths of the central barrier on one side of the screen, always opposite a 1 unit long barrier on the other side) and Condition (Congruent vs. Incongruent) as factors (Figure 2.3a). This analysis yielded a statistically significant main effect of Cost Asymmetry ($F(3, 69) = 4.83$, $p = .004$, $\eta^2 = .17$). Post-hoc tests revealed that this effect was due to a decrease in the proportion of efficient choices in trials with a 0.75 unit long barrier compared to shorter lengths (a post-hoc Bonferroni-corrected t-test comparing 0.75 with 0 found a statistically significant difference in the proportion of efficient choices, $t(23) = 3.20$, $p = .024$, $d = 0.65$, 95% CI = [0.04, 0.20]; whereas comparisons to 0.25 and 0.5 unit lengths, respectively, found only tendencies for higher efficiency ratios in trials with the shorter barriers: $t(23) = 2.80$, $p = .062$, $d = 0.57$, 95% CI = [0.02, 0.15]; $t(23) = 2.78$, $p = .064$, $d = 0.57$, 95% CI = [0.03, 0.20]). Neither the main effect of Condition ($F(1, 23) = 3.46$, $p = .076$, $\eta^2 = .13$), nor the interaction between Cost Asymmetry and Condition was statistically significant ($F(3, 69) = 1.48$, $p = .227$, $\eta^2 = .06$).

2.3.2 The effects of choices on performance

To test whether choosing the efficient path improved the accuracy and the speed of object transfer, mean frequencies of collisions per trial and mean trial duration were analyzed according to the decisions made. Considering that participants exhibited a strong tendency to make efficient choices throughout the experiment, the number of sub-efficient choices was much lower than that of efficient ones. Five participants did not make any sub-efficient choices. On average, participants completed the trials with the same level of accuracy when making efficient decisions ($n = 24$, $M = 0.10$, $SD = 0.07$) as when choosing the sub-efficient path ($n = 19$, $M = 0.15$, $SD = 0.20$; $t(18) = 1.10$, $p = .286$, $d = 0.25$, 95% confidence interval (CI) for the mean difference = [-0.05, 0.15]). However, a paired-samples t-test on mean trial durations demonstrated that the participants completed the task more slowly when making sub-efficient ($n = 19$, $M = 7.88$ s, $SD = 1.83$) than efficient choices ($n = 24$, $M = 6.06$ s, $SD = 1.02$; $t(18) = 9.72$, $p < .001$, $d = 2.23$, 95% CI of difference on log-transformed data = [0.09, 0.14]).

2.3.3 Exploratory analyses

To address the question whether participants had become more efficient over the course of the task, we compared the efficient choice ratios in the first half (Block 1, 40 trials) with the second half (Block 2, 40 trials) of the experiment using Wilcoxon signed-rank tests. We found that, in the Incongruent condition, the ratio of efficient choices increased between Blocks 1 and 2 (from M_{Block1} = 0.75, SD_{Block1} = 0.33, to M_{Block2} = 0.85, SD_{Block2} = 0.24, V = 5, p < .001, r = -.97, 95% CI = [-0.32, -0.13]). We observed no such increase in the Congruent condition (from M_{Block1} = 0.86, SD_{Block1} = 0.24, to M_{Block2} = 0.90, SD_{Block2} = 0.19, V = 36, p = .315, r = -.76, 95% CI = [-0.27, 0.07]). However, one-sample comparisons to chance also suggested that in Block 1, the ratios of efficient path choices were already significantly higher than chance level, regardless of condition (Congruent: V = 289, p < .001, r = .92, 95% confidence interval (CI) for the efficient choice proportion difference from chance [arcsine transformed chance level of 0.5 = 0.7854] = [0.36, 0.79]; Incongruent: V = 237, p = .013, r = .58, 95% CI = [0.15, 0.58]). In short, we found that the participants of Experiment 1 made efficient choices already in the first half of the experiment. However, on trials where taking the longer sub-path first was the efficient decision (Incongruent condition), participants chose it more frequently over time, suggesting that practice had some effect on making efficient choices.

2.3.4 Discussion

The participants acted efficiently, predominantly choosing the shorter total path to transfer the object. This was more pronounced for layouts where the difference in path length was larger, resulting in higher cost asymmetry. Choosing the shorter path resulted in shorter trial completion times. The tendency to choose the gap closer to the starting position in Neutral trials indicates that participants may have prioritized completing the first sub-goal (cf. Rosenbaum et al., 2014).

2.4 Experiment 2: Co-efficiency - Methods

To test the hypothesis that people maximize the co-efficiency of joint actions, pairs of participants performed the task together as a sequentially distributed joint action. The co-efficiency hypothesis predicts that the actor initiating the joint action should choose the sub-path that results in the shortest path for the dyad, rather than minimizing her own or her partner's movement distance.

2.4.1 Participants

Target sample size was determined in the same way as for Experiment 1, by conducting a power analysis in G*Power 3 (Faul et al., 2007). A sample size of $N = 24$ was estimated to provide 80% statistical power. Twenty-eight right-handed participants took part in Experiment 2. We excluded two pairs from data analysis due to a computing error ($n = 1$) and failure to understand the instructions ($n = 1$). We report the results of 12 dyads (4 mixed-gender and 4 female dyads; $N = 24$, 12 males, age $M = 25.4$ years, $SD = 4.14$). In all joint experiments (including SOM-R Experiments 3-4), we excluded dyads' data if they had previously known each other, to prevent any confound related to familiarity. For Experiment 2, we did not happen to recruit participant pairs who were familiar with each other.

2.4.2 Apparatus

We used the same apparatus as in Experiment 1. Stimulus presentation and data recording were controlled by a script of the task adapted for dyads, using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997; Kleiner et al., 2007) in MATLAB® (The MathWorks, Natick, MA). Two response boxes (Black Box Toolkit Ltd.) were used, one for each participant.

2.4.3 Stimuli and task

Experiment 2 employed the same stimuli and task as Experiment 1, with the difference that both members of the dyad had to act jointly to transfer the ball from the starting to the goal location: one participant moved the ball to the sub-goal location (i.e., one of the two gaps in the wall), and the other moved it from there to the goal location. Participants took turns in completing each part of the action sequence in a trial. The sub-goal of transporting the football to a gap in the middle of the screen was assigned to the decision-making participant, who acted first on the given trial (Actor 1). After Actor 1 handed over the football to their partner (Actor 2), she moved it from the gap to the goal location, and thus completed the task. The role of Actor 1 was randomly assigned throughout the task in each trial, and both participants acted as Actor 1 and Actor 2 an equal number of times.

2.4.4 Design

This was the same as in Experiment 1. The primary dependent variable was Actor 1's choice of sub-path to a sub-goal, i.e., to the gap where she would transfer the football to her partner, Actor 2. Accordingly, the main factor we manipulated was whether choosing the gap that offered the shorter sub-path to achieve Actor 1's goal of passing the ball to her partner resulted in a shorter total path for the dyad. When the central barrier was longer on Actor 1's side than the one on the other side (Figure 2.1a), maximizing either individual efficiency or co-efficiency required Actor 1 to choose the closer gap (Congruent trials). When the central barrier was longer on the opposite side (Figure 2.1b), maximizing co-efficiency required Actor 1 to opt for the farther gap, while maximizing her individual efficiency meant choosing the closer gap (Incongruent trials). When the barrier lengths on the two sides were equal (Figure 2.1c), either choice resulted in the same total path length (Neutral trials).

Congruent and Incongruent trials had the same levels of asymmetry between path lengths as in Experiment 1 (see different barrier lengths of the factor Cost Asymmetry). The list of trials from Experiment 1 was duplicated so that each participant completed the 80 trials used in Experiment 1. Trial order was fully randomized.

2.4.5 Procedure

Participants faced one another standing on the two opposite sides of the horizontally placed touchscreen and had full visual access to what their partner was doing (Figure 2.1d). Since we used a turn-taking task, only the acting player was in control of the ball. In the meantime, their partner had to keep a key pressed on the response box in front of them. Participants were instructed to finish each trial as accurately as possible while minimizing collisions, and to avoid communicating with one another during the task. The instructions also emphasized the shared goal of moving together the ball from one side of the screen to the other. Participants first completed a brief practice session of 10 trials, followed by the main experimental task. Finally, they filled in a short questionnaire on what they thought to be the experiment's purpose and how much they liked their partner, using a 7-point Likert scale (1 – Not at all, 7 – Very much).

At the beginning of each trial, when both actors pressed and held down the keys on their respective response boxes, they saw the layout of the game on-screen, which displayed their starting squares without the ball image. After 1500 ms, the football appeared in one of the squares. The actor with the object on their side (Actor 1) moved first and chose a transfer point to pass the ball over to her partner through one of the two circled gaps between the walls (Figure 2.1a-c). When the ball was fully inside the circle, the background of the circle was highlighted, any further movement of the ball by Actor 1 was blocked, and she had to press her response key again. Actor 2 then moved her hand from her respective response box key to the ball and dragged it back to the goal location on her side.

Two movement trajectories were registered: Actor 1's move to the gap from the starting location, and Actor 2's move from the gap to the goal location. A trial was complete when Actor 2 took the ball back to her home square (the goal location). No feedback was provided about speed or accuracy of performance. Each dyad completed the task in their own time. Participants completed the task on average in M = 21.49 minutes (SD = 2.89).

2.4.6 Data analysis

Data transformations and analyses were identical to those in Experiment 1. The primary dependent measure was the proportion of Actor 1's co-efficient choices, i.e., the shorter sub-path in the Congruent, and the longer sub-path in the Incongruent condition.

2.5 Results

2.5.1 Proportion of co-efficient choices

Participants opted for sub-paths that maximized the co-efficiency of the dyad (Figure 2.2b): one-sample Wilcoxon-tests indicated that in Congruent trials, participants passed over the ball through the gap closer to them significantly more often than chance (M = .85, SD = 0.14, V = 300, p < .001, r = 1.00, 95% confidence interval (CI) for the difference between the proportion of co-efficient choices and chance level [arcsine transformed chance level = 0.7854] = [1.13, 1.36]), whilst in Incongruent trials, they chose the gap farther away (M = .97, SD = 0.04, V = 300, p < .001, r = 1.00, 95% CI = [1.39, 1.48]). In Neutral trials, participants were significantly more likely to choose the longer sub-path on their side than the shorter one (M = .79, SD = 0.23, V = 277, p < .001, r = .85, 95% CI = [1.05, 1.31]). Paired-samples comparisons confirmed that the proportions of co-efficient choices were higher in both the Congruent trials (V = 300, p < .001, r = 1.00, 95% CI = [0.68, 1.00]) and in the Incongruent trials (V = 210, p < .001, r = .40, 95% CI = [0.25, 0.42]) than the proportions of the short and long sub-path choices in the Neutral trials, respectively. Furthermore, we found that three participants never chose sub-paths that were sub-efficient from the dyad's perspective.

A paired-samples comparison between the ratio of short sub-path choices in Congruent trials and long sub-path choices in Incongruent trials found that the ratio of co-efficient choices in the Incongruent trials was significantly higher than in the Congruent trials (V = 173.5, p = .002, r = .16, 95% CI = [0.13, 0.39], see Figure 2.2b). Participants made more co-efficient path choices when this meant reducing the effort of their partner than otherwise.

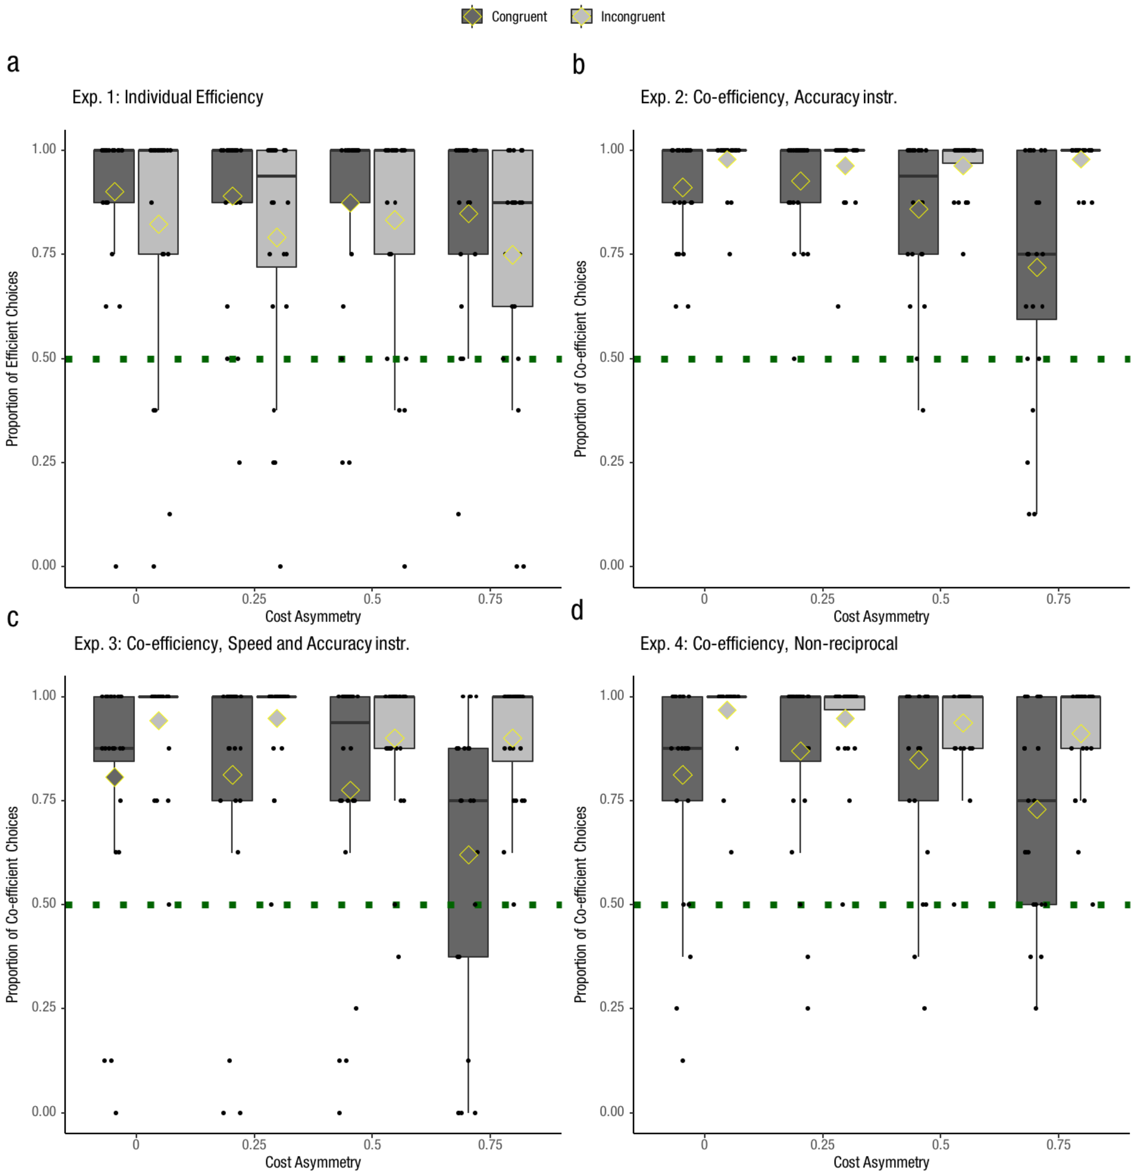

Figure 2.3: Illustrations of raw proportions of efficient path choices (N = 24) in the Congruent and Incongruent conditions, compared between layouts with different central barrier lengths in Experiments (a) 1, (b) 2, and in additional joint Experiments (c) 3 and (d) 4. Details and results of Experiments 3 and 4 may be found in the SOM-R. Each black dot represents a participant's efficient choice ratio in the given condition. Box plot lines indicate medians, diamonds indicate mean efficient choice proportions. The dotted lines indicate chance level (0.5) of choice proportion.

Efficient decisions were compared between different degrees of cost asymmetry in a 4 X 2 (Cost Asymmetry X Condition) repeated measures ANOVA on the ratios of short and long co-efficient sub-path choices. We found that the participants chose the co-efficient paths more often in Incongruent than in Congruent trials (main effect of Condition, $F(1, 23) = 17.13$, $p < .001$, $\eta^2 = .43$, Figure 2.3b). The participants chose co-efficient paths more frequently in trials with layouts with shorter barriers than in ones with longer barriers, as suggested by a statistically significant main effect of Cost Asymmetry ($F(3, 69) = 6.30$, $p < .001$, $\eta^2 = .22$). Furthermore, we found a statistically significant Cost Asymmetry X Condition interaction ($F(3, 69) = 7.48$, $p < .001$, $\eta^2 = .25$). This was due to a difference between the size of the Condition effect on proportions of co-efficient choices in trials with different degrees of Cost Asymmetry. Post-hoc Bonferroni-corrected paired-samples t-tests yielded statistically significant effects of Condition on the ratio of co-efficient choices in trials with 0, 0.5 and 0.75 unit long barriers, respectively (0 unit: $t(23) = 2.73$, $p = .048$, $d = 0.56$, 95% CI = [0.04, 0.25]; 0.5 unit: $t(23) = 2.98$, $p = .028$, $d = 0.61$, 95% CI = [0.05, 0.30]; 0.75 unit: $t(23) = 4.74$, $p < .001$, $d = 0.97$, 95% CI = [0.23, 0.59]), but not in trials with 0.25 unit long barriers ($p = 1.000$). We found that for most combinations of barrier lengths, it was true that Actor 1 made more co-efficient decisions when co-efficiency entailed helping her partner by choosing the gap farther away from herself (Incongruent trials), rather than choosing the gap closer to her (Congruent trials).

2.5.2 The effect of choices on performance

To test whether Actor 1's co-efficient choices improved the dyad's performance, we compared mean numbers of collisions per trial and mean trial duration between trials where Actor 1 chose the co-efficient sub-path and those where she chose the sub-efficient sub-path. On average, dyads completed trials with a significantly higher level of accuracy when Actor 1 chose the co-efficient sub-path, colliding with onscreen walls fewer times ($n = 24$, $M = 0.16$, $SD = 0.10$) than when she chose the sub-efficient path ($n = 21$, $M = 0.33$, $SD = 0.42$, $t(20) = 2.18$, $p = .041$, $d = 0.48$, 95% CI = [0.01, 0.34]). Although actors were not explicitly instructed to optimize speed, making co-efficient decisions also resulted in shorter trial completion times. Trial duration was significantly longer for sub-efficient choices ($n = 21$, $M = 10.7$ s, $SD = 5.62$) than for efficient choices ($n = 24$, $M = 7.55$ s, $SD = 1.02$, $t(20) = 5.85$, $p < .001$, $d = 1.28$, 95% CI = [0.08, 0.17]).

2.5.3 Questionnaires

In the questionnaire addressing the perceived purpose of our study, a third of the participants mentioned that they considered the experiment as investigating cooperation ($n = 7$) and helping tendencies (n = 8). A minority made explicit reference to rational decision-making or optimization (n = 4), finding the shortest path for both players (n = 5), reactivity to a partner's actions (n = 6), and a few people thought we looked at the effect of getting tired, or being good at perceiving visual differences in distances (n = 3).

The rating of partners were generally high (Mdn = 6, SD = 0.95). The correlation (Spearman's ϱ) between liking ratings and the arcsine transformed ratios of co-efficient choices was not different from zero in either condition (Congruent: ϱ = .321, p = .126, Incongruent: ϱ = -.076, p = .725).

2.5.4 Exploratory analyses

Similarly to Experiment 1, we conducted additional exploratory analyses to address the potential influence of practice on efficient decision-making by comparing co-efficient choice ratios between the first half (Block 1, 80 trials) and the second half (Block 2, 80 trials) of the joint task. Paired-samples Wilcoxon signed-rank tests found that in the Congruent condition, the ratio of co-efficient choices increased between Blocks 1 and 2 (from M_{Block1} = 0.82, SD_{Block1} = 0.19, to M_{Block2} = 0.89, SD_{Block2} = 0.13, V = 35.5, p = .031, r = -.76, 95% CI = [-0.28, -0.01]). No such increase was observed in the Incongruent condition (from M_{Block1} = 0.96, SD_{Block1} = 0.10, to M_{Block2} = 0.98, SD_{Block2} = 0.04, V = 35, p = .484, r = -.77, 95% CI = [-0.31, 0.24]). Proportions of co-efficient choices were already significantly higher than chance in Block 1, regardless of condition (all ps < .001).

To investigate potential between-experiment differences in the ratios of efficient (Experiment 1) and co-efficient (Experiment 2) choices, we compared the ratios of (co-)efficient decisions in the Congruent and Incongruent conditions separately. Mann-Whitney U tests with Experiment as a factor found no statistically significant difference in the ratio of (co-)efficient choices in the Congruent condition (Experiment 1: M = 0.88, SD = 0.21; Experiment 2: M = 0.85, SD = 0.14; U = 351.5, p = .184, r = .22, 95% CI for the median difference between the two experiments = [-3.49e-5, 0.31]). In contrast, in the Incongruent condition, dyads in Experiment 2 made a statistically significantly higher proportion of co-efficient choices than individuals in

Experiment 1 made efficient choices (Experiment 1: $M = 0.80$, $SD = 0.28$; Experiment 2: $M = 0.97$, $SD = 0.04$; $U = 168$, $p = .011$, $r = -.42$, 95% CI: [-0.31, -9.572^{e-6}]). This asymmetric pattern in between-experimental differences suggests that facilitating a partner's actions in the joint task by taking the longer sub-path might have further boosted the ratio of co-efficient choices.

2.5.5 Discussion

When participants had multiple options to plan a movement in a coordination context, they considered not just their own, but also their partner's costs. This was demonstrated by the first actors' strong tendency to choose the sub-path that was more co-efficient, whether it resulted in reducing or increasing their partner's costs. That is, action initiators went for the shorter sub-path for themselves and the longer one for the partner in the Congruent condition and displayed the opposite pattern of choices in the Incongruent condition. When co-efficiency was unaffected by sub-path choices (Neutral trials), participants reduced their partner's costs.

2.6 General Discussion

Our experiments addressed the question whether people minimize the aggregate costs of actions when cooperating with others to reach a shared goal. We operationalized action costs as path length travelled while moving an object. We found that actors chose to minimize the total path length when offered two path options to complete a movement sequence. In the joint task, these total paths were distributed over co-actors, suggesting that participants aimed at maximizing the co-efficiency of the dyad. In the individual task, the choices were similar to joint performance, demonstrating efficient planning for the entire action sequence.

The decisions in the dyadic Incongruent condition, where taking a longer sub-path to a gap was analogous to reducing the partner's effort in joint object manipulation tasks (Meyer et al., 2013; Dötsch & Schubö, 2015), indicated that actors integrated their partners' effort into their planning and were motivated to reduce their partner's costs. However, in the Congruent condition participants refrained from reducing their partner's effort, maximizing the group's efficiency by forcing the partner to move along the longer sub-path. The complementary pattern of the two conditions suggests that, in joint action contexts, people aim at reducing aggregate group costs rather than minimizing the effort of either party. This is in line with Santamaria and Rosenbaum's shared-effort model (2011), which postulated that actors coordinate their actions to reduce aggregate costs of a group.

Two additional experiments tested the robustness of co-efficiency maximization (see the 2.7 SOM-R: Additional experiments - Methods). The results of Experiment 2 were replicated when participants were instructed to complete the task as quickly as possible in addition to being accurate (Experiment 3), and when the identity of the decision-maker was fixed to eliminate turn-taking of choices (Experiment 4). The latter results indicated that expectation of reciprocity is not necessary for efficient joint action planning.

Notably, the Congruent and Incongruent trials induced similar decisions already in the first half of the task in both experiments. This raises the possibility that, in the joint task, decision-makers disregarded their partners entirely when planning their actions and considered only the total path options that they could have executed individually. The differential results of the Neutral trials, however, provide evidence against this account: when co-efficiency did not discriminate between the options, participants reduced their partner's costs by covering the longer distance (Experiment 2) but were biased in the opposite direction when they acted alone (Experiment 1). Furthermore, in Experiment 1, participants maximized efficiency similarly across conditions, whereas in Experiment 2, they made more co-efficient decisions in the Incongruent than in the Congruent condition. Lastly, we observed higher proportion of co-efficient choices in the Incongruent condition of Experiment 2 relative to Experiment 1. In other words, actors sacrificed the efficiency of their initial act more when this choice reduced the partner's effort than when it increased the partner's costs, or when they performed the task alone. These findings suggest that the participants planned the joint action sequences with their partners in mind, possibly even signaling cooperative attitudes by taking over effort from them when this decision did not compromise co-efficiency.

Future experiments should address the mechanism underlying co-efficiency maximization in more detail. Candidate mechanisms for such decision-making include a rational calculus of joint costs, which sums agent-specific individual costs, along with the use of heuristics, such as simulating entire action sequences to be performed by the individual alone. Beyond specifying the mechanism, a model of rational joint action planning will need to explore the boundary conditions of co-efficiency maximization. In the present study we focused on path length, but actions may similarly be optimized for exerted effort, in which case movement curvature could also be considered. Finally, joint optimization could be modulated by benefit sharing, and asymmetries in competence or in access to information.

2.7 SOM-R: Additional experiments - Methods

In Experiment 3, we tested the co-efficiency maximization hypothesis under different task instructions. Experiment 4 addressed the hypothesis that co-efficient decisions may have been the result of a direct reciprocity strategy of a partner's efforts: participants performed a non-reciprocal version of the joint task.

2.7.1 Participants

Thirty right-handed participants took part in Experiment 3, forming dyads. We excluded three pairs from data analysis because the participants had previously known one another. We report the results of 12 dyads (9 mixed-gender and 3 females; *N* = 24, 9 males, *M* = 24.5 years, *SD* = 2.59).

Fifty-four (27 pairs of) right-handed participants took part in Experiment 4. Three pairs were excluded due to misunderstanding instructions, experimenter error and because the co-actors knew each other, respectively. Results for 24 actors from 24 dyads (10 mixed-gender, 12 females; *N* = 24, 9 males, *M* = 23 years, *SD* = 3.41) are reported.

2.7.2 Apparatus

We used the same apparatus as in Experiment 2.

2.7.3 Task, Design and Procedure

The task, design and procedure in Experiments 3 and 4 were identical to Experiment 2, bar the following modifications: in Experiment 3, participants were given a different instruction (i.e., to finish each round as quickly and accurately as possible); in Experiment 4, only one person from each dyad made decisions as Actor 1. In Experiment 3, each participant completed 80 trials, totaling 160 trials for a dyad. In Experiment 4, each session comprised only 80 trials as only one person started the trials in each dyad. Participants in Experiment 3 completed the task on average in *M* = 16.01 minutes (*SD* = 6.84) and in Experiment 4, in *M* = 14.91 minutes (*SD* = 2.62), while in Experiment 2, dyads took on average 21.49 mins. In Experiment 3, we instructed participants to be fast in addition to being accurate, which is the reason why they were faster than in Experiment 2. In Experiment 4, the instruction only mentioned accuracy, just like in Experiment 2. Here, one potential explanation for being overall faster may be that in Experiment 4, only one person decided about distributing workload between co-actors. This might have reduced degrees of freedom for overall decision-making, which could have manifested itself in faster completion of the experiment.

2.7.4 Data analyses

Data transformations and analyses were identical to Experiment 2's.

2.8 SOM-R: Results

In both experiments, Actor 1 chose sub-paths that maximized the dyad's co-efficiency (Figure 2.2c and d in the main text). On Congruent trials, actors mostly passed the football over to their partner in the gap closer to themselves (Experiment 3: M = 0.75, SD = 0.30, V = 256, p = .003, r = .71, 95% confidence interval (CI) for the co-efficient choice proportion difference from chance level [arcsine transformed chance level of 0.5 = 0.7854] = [0.97, 1.29]; Experiment 4: M = 0.82, SD = 0.22, V = 287, p < .001, r = .91, 95% CI = [1.07, 1.39]); on Incongruent trials, they chose the more distant gap (Experiment 3: M = 0.92, SD = 0.12, V = 300, p < .001, r = 1.00, 95% CI = [1.28, 1.48]; Experiment 4: M = 0.94, SD = 0.10, V = 300, p < .001, r = 1.00, 95% CI = [1.32, 1.48]). Furthermore, participants made more co-efficient decisions when this entailed facilitating their partner's action by taking the longer sub-path themselves, than when it meant taking the short sub-path (Experiment 3: V = 219, p = .014, r = .46, 95% CI = [0.05, 0.47]; Experiment 4: V = 172, p = .002, r = .15, 95% CI = [0.12, 0.37]). Finally, we also replicated the altruistic bias on Neutral trials: Actor 1 chose the longer sub-path significantly more often than expected by chance (Experiment 3: M = 0.63, SD = 0.28, V = 204.5, p = .044, r = .36, 95% CI = [0.79, 1.10]; Experiment 4: M = 0.63, SD = 0.21, V = 217.5, p = .016, r = .45, 95% CI = [0.82, 1.06]).

To investigate the effect of cost asymmetries between partners on decision-making, we compared co-efficient sub-path choice proportions between each barrier length (Figure 2.3c and d in the main text). In both experiments, 4 X 2 (Cost Asymmetry X Condition) repeated-measures ANOVAs revealed statistically significant main effects for Cost Asymmetry (Experiment 3: $F(3, 69) = 13.36$, $p < .001$, $\eta^2 = .37$; Experiment 4: Greenhouse-Geisser corrected $F(2.17, 49.86) = 11.00$, $p < .001$, $\eta^2 = .32$) and Condition (Experiment 3: $F(1, 23) = 6.78$, $p = .016$, $\eta^2 = .23$; Experiment 4: $F(1, 23) = 13.69$, $p = .001$, $\eta^2 = .37$), and interactions between Cost Asymmetry and Condition (Experiment 3: Greenhouse-Geisser corrected $F(2.11, 48.48) = 4.00$, $p = .023$, $\eta^2 = .15$; Experiment 4: Greenhouse-Geisser corrected $F(2.07, 47.50) = 3.80$, $p = .028$, $\eta^2 = .14$). Post-hoc Bonferroni-corrected paired-samples t-tests yielded statistically significant effects of Condition on co-efficient choice proportions in trials with 0.75 unit length barriers (Experiment 3: $t(23) = 3.31$, $p = .012$, $d = 0.68$, 95% CI = [0.17, 0.73]; Experiment 4: $t(23) = 3.82$, $p = .004$, $d = 0.78$, 95% CI = [0.13, 0.43]), and in Experiment 4, also in trials with 0 unit length barriers ($t(23) = 3.74$, $p = .004$, $d = 0.76$, 95% CI = [0.12, 0.42]). Participants made more co-efficient choices in Incongruent than in Congruent trials.

Co-efficient decisions led to faster and more accurate performance (results of paired-samples t-tests are reported in Table 2.1).

To explore the effect of practice, we compared co-efficient choice ratios between the first and the second half of each experiment (Block 1 v. Block 2, Experiment 3: 80 trials/block, Experiment 4: 40 trials/block), with paired-samples Wilcoxon signed-rank tests. In both experiments, the ratios of co-efficient choices were significantly higher than chance already in Block 1, regardless of condition (all ps < .05). There was a statistically significant increase in the ratio of co-efficient choices between Blocks 1 and 2 in the Congruent condition (Experiment 3: $M_{\text{Block1}} = 0.72$, $SD_{\text{Block1}} = 0.30$ to $M_{\text{Block2}} = 0.80$, $SD_{\text{Block2}} = 0.31$, $V = 38$, $p = .007$, $r = -.75$, 95% CI = [-0.28, -0.07]; Experiment 4: $M_{\text{Block1}} = 0.76$, $SD_{\text{Block1}} = 0.27$ to $M_{\text{Block2}} = 0.88$, $SD_{\text{Block2}} = 0.19$, $V = 16$, $p = .005$, $r = -.89$, 95% CI = [-0.47, -0.12]), as participants selected the co-efficient short path more often in the second than in the first half of the experiment. In Experiment 3, but not in Experiment 4, we also found a significant increase in the ratio of co-efficient choices in the Incongruent condition (Experiment 3: $M_{\text{Block1}} = 0.88$, $SD_{\text{Block1}} = 0.21$ to $M_{\text{Block2}} = 0.96$, $SD_{\text{Block2}} = 0.10$, $V = 16$, $p = .042$, $r = -.89$, 95% CI = [-0.54, -0.01]; Experiment 4: $M_{\text{Block1}} = 0.93$, $SD_{\text{Block1}} = 0.13$ to $M_{\text{Block2}} = 0.96$, $SD_{\text{Block2}} = 0.08$; $V = 22.5$, $p = .373$, $r = -.85$, 95% CI = [-0.35, 0.06]).

Table 2.1: Comparisons of speed and accuracy measures between trials where participants made co-efficient and sub-efficient choices

	Co-efficient trials	**Sub-efficient trials**	**Statistic (t)**	***df***	***p***	***d***	**95% CI for mean difference**
Experiment 3							
Collision	$n = 24$ $M = 0.28$ $SD = 0.11$	$n = 24$ $M = 0.65$ $SD = 0.53$	3.34	23	.003*	0.68	[0.14, 0.59]
Trial duration (s)	$n = 24$ $M = 5.60$ $SD = 0.67$	$n = 24$ $M = 7.17$ $SD = 1.27$	8.94	23	< .001**	1.82	[0.08, 0.13]
Experiment 4							
Collision	$n = 24$ $M = 0.11$ $SD = 0.06$	$n = 19$ $M = 0.36$ $SD = 0.48$	2.31	18	.033*	0.53	[0.02, 0.48]
Trial duration (s)	$n = 24$ $M = 7.46$ $SD = 1.36$	$n = 19$ $M = 9.68$ $SD = 2.17$	11.01	18	< .001**	2.53	[0.09, 0.13]

*p < .05. **p < .01.

Finally, to see whether dyads acted more efficiently than individuals, we compared the ratio of co-efficient choices in Experiments 3 and 4 to the ratio of efficient choices in Experiment 1 (Individual task, see main text). Mann-Whitney U tests with Experiment as a factor found that the ratio of efficient choices in the Congruent condition was statistically significantly higher in

Experiment 1 than in Experiment 3 ($U = 398.5$, $p = .021$, $r = .38$, 95% CI for the median difference between the two experiments = [0.00, 0.41]), but the difference was not statistically significantly different between Experiments 1 and 4 ($p = .179$). In the case of the Incongruent conditions, we found that dyads in Experiment 4 chose the co-efficient long sub-paths significantly more often than individuals chose the efficient paths in Experiment 1 ($U = 194.5$, $p = .047$, $r = -.33$, 95% CI = [-0.29, 0.00]), although this difference did not reach statistical significance between Experiments 1 and 3 ($p = .100$). In sum, we found no consistent evidence that dyads made overall more (co-) efficient decisions than the individual participants of Experiment 1.

We calculated the correlation (Spearman's ϱ) between liking ratings (Experiment 3: $Mdn = 6$, $SD = 1.28$; Experiment 4: $Mdn = 6$, SD = 1.26) and the arcsine transformed ratios of co-efficient choices, which was not different from zero in either condition (Congruent: Experiment 3: $\varrho = -.119$, $p = .587$, Experiment 4: $\varrho = -.303$, $p = .151$; Incongruent: Experiment 3: $\varrho = .252$, $p = .246$, Experiment 4: $\varrho = -.346$, $p = .097$).

2.9 SOM-R: Discussion

In Experiment 3, in the experimental conditions, dyads optimized their actions by making decisions that minimized joint costs already from the first half of the task, and in the Neutral condition, we replicated the bias to facilitate a partner's action. We therefore replicated our results under different task instructions, highlighting both the accuracy and speed of performance.

In Experiment 4 we replicated all results of previous joint experiments, indicating that co-efficient decisions are not driven by reciprocity expectations.

Chapter 3. Computing joint action costs: Co-actors minimize the aggregate individual costs in an action sequence

3.1 Introduction

Humans cooperate by sharing goals with others, and by planning and coordinating actions with their partners to achieve these goals (Bratman, 1992; Butterfill, 2016; Sebanz, Bekkering et al., 2006). Everyday social interactions, such as assembling furniture with a friend, or cooking a meal together, attest to the complexity of planning required in cooperative activities. For instance, family members might share the overarching goal of cooking a paella: each of them represents and works towards specific sub-goals (e.g., chopping vegetables), and they need to distribute the necessary actions among themselves. Many sub-tasks contribute to the joint goal of cooking a paella. Accordingly, these actions may be executed in many different ways, with varying degrees of efficiency. When people distribute sub-tasks between themselves, the individual efficiencies of co-actors often depend on each other; in some situations, they are inversely related. That is, because a complex joint action may be composed of many interdependent sub-tasks, performing a less effortful sub-task may force the other person to contribute a more effortful complementary action to ensure the success of the joint action.

Planning cooperative sequences can be regarded as making a series of decisions about actions to be performed by co-actors (cf. Wolpert & Landy, 2012, on individual motor planning). What principles guide people's decision-making? 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 et al., 2016; Santamaria & Rosenbaum, 2011; Török et al., 2019). These findings are paralleled by evidence that in certain economic games, people sometimes make decisions consistent with a collective utility-maximizing strategy based on team reasoning (Sugden, 2003), rather than choosing individual utility-maximizing solutions

(Colman et al., 2008a). We argue that such behaviors do not only appear in contexts with financial rewards at stake.

In the joint action task of Török and colleagues (2019) participants made binary decisions between two action plans to coordinate their hand movements with a partner in a sequential manner. One of the two options was more efficient for the decision-making actor (i.e., the initiator of the action sequence), the other option was more efficient for her partner, therefore the jointly efficient plan was more individually efficient for either the decision-maker or for the partner. The participants tended to make *co-efficient* (i.e., jointly efficient) rather than individually efficient decisions that would have either maximized personal efficiency or would have altruistically increased the utility of a partner (cf. Axelrod & Hamilton, 1981; Trivers, 1971). The present study investigated the computations behind decisions that minimize the aggregate costs of the group.

To minimize a dyad's costs in action planning, a decision-maker first needs to estimate them (Körding & Wolpert, 2006). In the case of joint actions, we assume that the expected individual costs of potential joint action sequences are integrated to achieve optimal decisions. People are sensitive to their real or virtual partner's individual efforts, needs and task difficulty (Chennells & Michael, 2018; Ray & Welsh, 2011; Ray et al., 2017). We hypothesize that, whenever this is calculable, the costs of joint actions are estimated as the summed total of individual costs. This proposal is compatible with computational work in which joint utility is 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 a generic process to achieve co-efficiency (applicable when we conceptualize the joint-cost estimation problem in the abstract, akin to a mathematical problem of combining two quantities), depending on the actual context, shortcuts may also be available. For example, Török and colleagues' (2019) task required two actors to move an object along one of two paths by taking turns. While the movement was divided between the participants, the decision-making actor could have 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 result in choosing the co-efficient action option from the alternatives without requiring the planner to sum individual costs.

In the present study, we employed a task in which joint action costs cannot be computed without representing and summing two individual action costs. The participants had to move objects on a touchscreen sequentially[7], and the cost of this action was assumed to be proportional to the path length of movement. Crucially, the physical separation of paths to be taken by co-actors 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 and joint action costs in various ways. Since a priori we could not exclude the possibility that participants might adapt to the correlational structure of the cost parameters during the experiment, we used separate participant samples for three versions of the task in which pairs of cost parameters were de-correlated. We observed highly consistent results across experiments, and thus we report the analyses of pooled data. The individual experiments' details are available in the Supplementary Material (section 3.5.2 Results – Additional Experiment-wise Information and Table 3.3).

If people represent the joint costs of an action sequence as a weighted 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 social hierarchy (Kleiman-Weiner et al., 2016), the individual costs of the actors would be weighed equally in the sum.

We also investigated an alternative hypothesis, according to which the equality of contribution matters more than the efficiency of joint performance. People are often motivated to reduce payoff inequality in economic games (Dawes et al., 2007; Fehr & Schmidt, 1999), and it is possible that in a joint action context they minimize the difference in the action costs distributed across co-actors rather than maximizing the expected utility of the dyad. Such decisions may be

[7] Despite its sequential nature, we consider this a joint action. For an action to qualify as joint action, a goal that is not individual but shared between co-actors should be present. In our task, the interdependence of the two individual actions to reach the goal of matching object pairs ensures that there is a joint goal.

based on a motivation to be fair to an interaction partner (Rand et al., 2013), although recent results support the co-efficiency hypothesis against the trial-based fairness account (Strachan & Török, 2020).

3.2 Methods

3.2.1 Participants

We recruited participants through our institution's Research Participation System and a student job agency. They gave their informed consent and received vouchers for their participation. The study was approved by the local ethics committee. To ensure that the present study was adequately powered to make inferences in the Bayesian model, the target sample size was set to 20 dyads (40 participants) per experiment, a larger sample size relative to Török and colleagues' (2019) study. We present the data of 120 participants (82 females, 2 preferred not to specify; M_{age} = 23.81 years, *SD* = 4.07) (see section 3.5.1 Methods – Additional Experiment-wise Information1 for descriptions of exclusions).

3.2.2 Apparatus

The task was performed on a touchscreen monitor (Iiyama PROLITE 46", resolution 1920 X 1080 pixels, separate sync - horizontal: 31.47 – 67.5 KHz, vertical: 47 – 63 Hz) lying flat on a table between two participants facing each other. Stimulus presentation and data recording were controlled by a script using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997; Kleiner et al., 2007) in MATLAB® (The MathWorks, Natick, MA). Two response boxes (Black Box Toolkit Ltd.) were used to control trial onset.

3.2.3 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 the two participants' task areas, (2) two pairs of black target objects (two circles and two squares, 30 px diameter) distributed between the task areas, and (3) two black-bordered octagonal starting locations (96 x 96 px) with another, small octagon inside

(60 x 60 px, Figure 3.1). The starting locations were always displayed at mid-position along the longer sides of the screen, aligned with the response box buttons.

Participants were instructed to keep their dominant index finger on the button of their response boxes to trigger the start of each trial. At the beginning of each trial, one of the smaller octagons was orange-colored to signal which participant would initiate the joint action as Actor 1. In each trial, Actor 1 had to choose between the two target objects on her side and drag it back to her starting location. The 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.

After 3 seconds, the color switched to green, which served as a cue for Actor 1 to start moving. By dragging the green octagon over a black object with her index finger, the participant picked up the object and collected it by transferring it back to her starting location. Once Actor 1 collected one of the objects, she pressed the button on her response box again to make the white octagon in front of her partner (Actor 2) turn green. The appearance of this second green octagon cued Actor 2 to start moving to collect the matching object in his task area. 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). Thus, while both participants acted in each trial, only Actor 1 made the decision that determined the individual and joint costs incurred during the completion of the task.

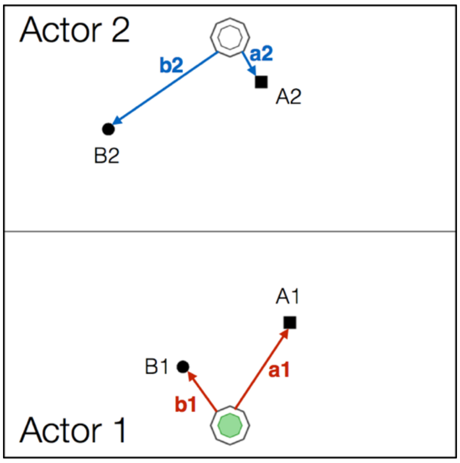

Figure 3.1: An example of the layouts presented to the participants. Starting locations were indicated by the octagons, and the locations of the two pairs of black target objects (circles and squares) were generated by stochastic selection processes. The arrows and labels (not shown to the participants) indicate the distances between the starting locations and the target objects, which provided the basis for cost calculations comparing the two target options.

3.2.4 Design

We considered the cost of an action as a monotonic function of the path length that the object covered on the touchscreen when dragged, and, for the sake of simplicity, we treated the path length as the absolute cost paid for its transport. For example, in Figure 3.1, 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 their difference: a1-b1. We call this value Self Cost Disparity, or simply Self Disparity. If Actor 1 intends to make individually efficient decisions, she should choose A1 as the target when the Self Cost Disparity is negative, and B1 when this value is positive (as is the case in Figure 3.1). 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 because 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 arrangement (Figure 3.1), the Joint Disparity is negative, suggesting that from the dyad's perspective, collecting the square objects 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 (negative Other Disparity), but not for Actor 1 (positive Self Disparity).

We assume that the likelihood of choosing object A1, which was always the square in the decision-maker's side of the screen, 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, forcing Actor 2 to act on A2.

To generate the target objects' locations, we first sampled Self Disparity and Other Disparity (section 3.5.1 Methods – Additional Experiment-wise Information2 reports details of the sampling procedure). The positions of the objects were then randomly selected to match these disparities. For each dyad, 100 different spatial arrangements were generated and repeated, once per each participant acting as Actor 1. Trial order was pseudo-randomized, with the constraint that neither of the participants be assigned the role of Actor 1 more than 3 times in a row.

To address our alternative hypothesis, we operationalized Fairness 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: the difference between [abs(a1-a2)] and [abs(b1-b2)] (Figure 3.1). Choices were considered 'fair' if the object Actor 1 picked up was associated with a relatively smaller difference in the path lengths between Actors 1 and 2 than the path length difference associated with the alternative object. When the fairness measure was negative, choosing object A1 was the fair option; when positive, object B1 was fair.

3.2.5 Procedure

Before the object matching task, the participants read step-by-step instructions on how to complete a trial. They were instructed to collect matching object pairs by cooperating with their partner, without communicating, and to complete each trial as quickly as possible. The participants first completed 6 practice trials to familiarize themselves with the task, the touchscreen and the response box buttons. They then completed the main task (on average in $M = 34.62$ minutes, $SD = 4.99$) without receiving any feedback.

3.2.6 Data Analysis

To test whether object choices were influenced primarily by the difference between joint action costs, rather than between the individual costs of Actor 1 or Actor 2, we fitted and contrasted three Bayesian hierarchical logistic regression models (Kruschke, 2015; details of the model are reported in 3.5.1 Methods – Additional Experiment-wise Information3). Specifically, the probability of 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*. We expected that the third model would have the best fit to the data, since it is the only one that can express the co-efficient strategy which dictates that actors should equally weigh their own and their partner's cost disparities. Additionally, we fitted models that predicted choices by (4) Fairness or by (5) the linear combination of Self and Other Disparities and Fairness.

The posterior distributions of the beta coefficients were estimated in JAGS (Plummer, 2003) with the *runjags* package in R (Denwood, 2016). We report the population-level estimates. To compare models, we calculated Leave-one-out cross-validation measures (LOO-CV; Gelman et al., 2014; Vehtari et al., 2017) using the *loo* package in R (Vehtari et al., 2020). The lower the information criterion for a model, the better its expected accuracy at out-of-sample prediction of future data. We also compared Area Under the Curve measures for each model (AUC; Fawcett, 2006), quantifying model fit to the observed data. We base model comparison on the AUC and LOO-CV (see Table 3.3).

3.3 Results

3.3.1 Descriptive Statistics

The participants chose object A1 on 49.7% of trials. The individual object choice proportions were not different from chance (Wilcoxon signed-rank test against 0.5: $V = 2724$, $p = .726$, rank-biserial correlation $r = -.25$, 95% confidence interval (CI) for proportion .50 = [.49, .51]).

3.3.1.1 Cost-minimization. The magnitudes of cost disparities strongly influenced object choices. Overall, participants chose the object resulting in a co-efficient action sequence 77% of the time (95% CI for proportion .77 = [.76, .79]), which was significantly higher than chance at 66%[8] ($V = 7260$, $p < .001$, $r = 1.00$). The participants' decisions are illustrated in Figure 3.2a, together with the predictions calculated for each cost-minimizing strategy (Figure 3.2d-f).

[8] In 66% of the trials, the predictions of Self- and Joint cost-minimization overlapped, so we used that as chance level.

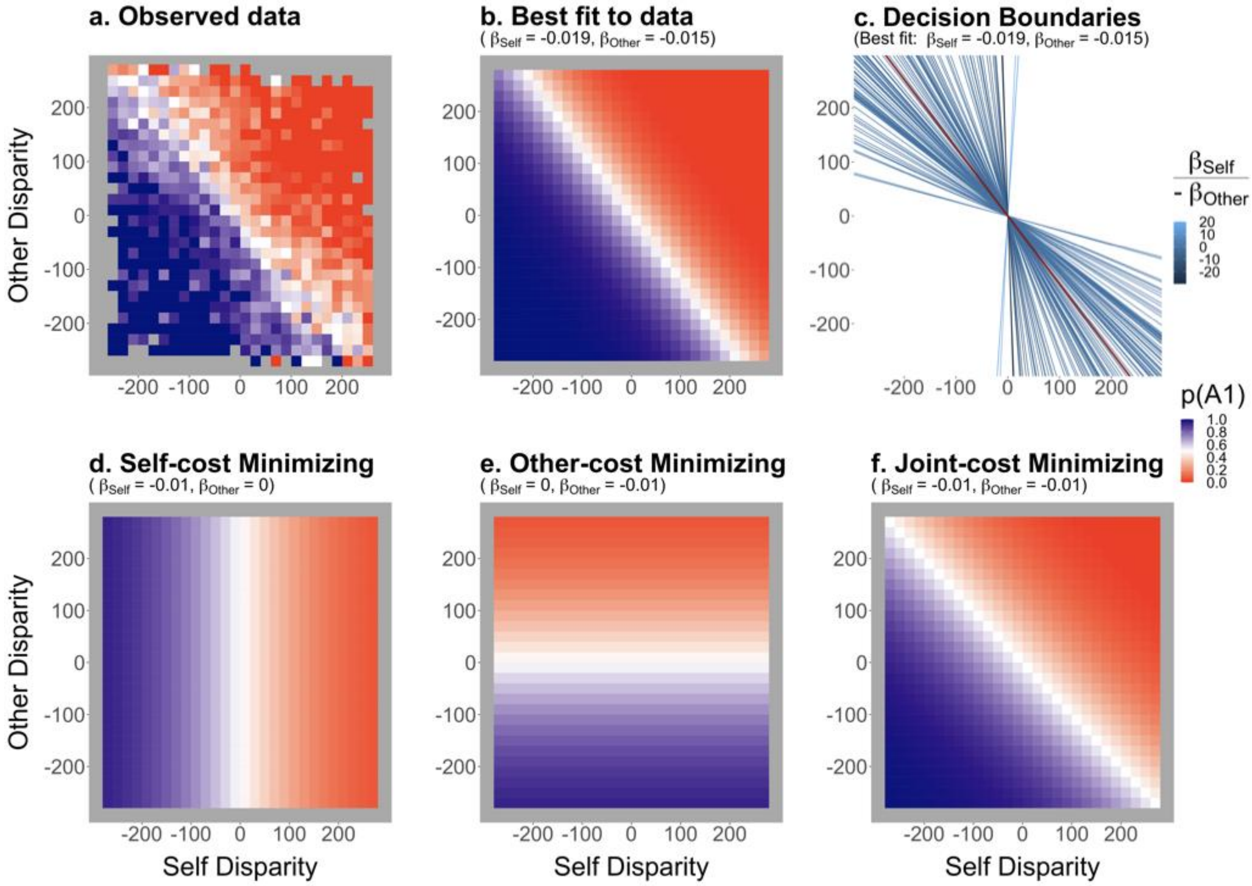

Figure 3.2: (a) Observed object A1 choices, and (b) the posterior predictions of the best-fitting model using the linear combination of Self and Other Disparities. (c) Individual decision boundaries according to the best-fitting model. The red line indicates the population-level boundary. (d-f) Predictions for optimal responses according to Self, Other, and Joint (i.e., Self + Other) cost-minimizing strategies, respectively. The lower the disparity to be minimized according to a model, the higher the probability of picking object A1 (blue). Predictions were calculated assuming that one pixel increase in a given parameter would result in 1% decrease in the odds of choosing A1 over B1. All plots feature disparities in pixels.

3.3.1.2 Fairness. The participants made fair choices 47.5% of the time (95% CI for proportion .48 = [.46, .50]), which was not significantly different from chance (V = 2479, p = .047[9], r = -.32). Decisions were more strongly influenced by Joint-Cost Minimizing concerns than by Fairness (Figure 3.3).

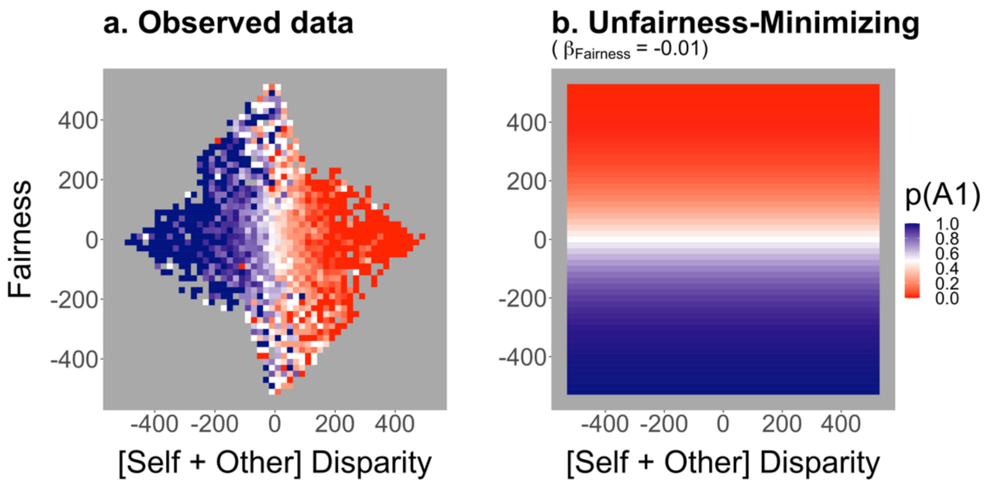

Figure 3.3: (a) Observed object A1 choices as a function of Fairness and the linear combination of Self and Other Disparity. (b) Predictions for optimal responses according to a strategy that minimizes the unfairness of task distribution between co-actors. The lower the degree of asymmetry in costs (or the magnitude of joint disparities) the higher the probability of picking object A1. The predictions were calculated assuming that one pixel increase in unfairness would result in 1% decrease in the odds of choosing A1 over B1. Both plots feature disparities in pixels.

3.3.2 Parameter Estimations

The posterior modes and 95% highest density intervals (HDI) for the population-level parameters that represent how participants weighted the cost disparities to make object choices are presented in Table 3.1 for all of the models fitted. Experiment-level parameter estimates are summarized in the SM (section 3.5.2 Results – Additional Experiment-wise Information , Table 3.3).

Turning first to the main hypothesis on the relative weighting of the Self and Other Disparities, a model including both disparities was a better fit for the data than models including either disparity alone. In this two-predictor model, both population-level means ($\mu_{\beta Self}$ and $\mu_{\beta Other}$) of the coefficients for the disparities were distributed below zero (Figure 3.4c, Self: 95% HDI: [-

[9] For three comparisons to chance, α-levels were Bonferroni-corrected for repeated testing (α = .017). As a measure of effect size, we report rank-biserial correlations (Kerby, 2014).

0.469, -0.273], Mode $\mu_{\beta Self}$ = -0.367; Other: 95% HDI: [-0.390, -0.193], Mode $\mu_{\beta Other}$ = -0.291). This suggests that both parameters made non-null contributions in the predicted directions. Increasing Self and Other disparities by a centimeter led to a 30.7% and 25.2% decrease in the odds of picking A1 over B1, respectively. The 95% HDI of the posterior of the difference between the disparities' coefficients included zero ($Mode_{diff(\mu_{\beta Self} - \mu_{\beta Other})}$ = -0.072, 95% HDI = [-0.214, 0.052]), suggesting that the difference between the magnitudes of the two disparities' effects on decision-making was credibly null. The average relative weights on Self and Other Disparity were .56 (95% HDI = [.42, .71]) and .44 (95% HDI = [.29, .59]), respectively (for participant-wise estimates, see Figure 3.9). This pattern of weights was not due to selfish and altruistic people's results averaging out: 95 participants' HDIs overlapped.

Second, Fairness on its own was not a meaningful predictor. When included in the model using Self and Other Disparities as predictors, the conditional influence of Fairness was still credibly null ($\mu_{\beta Fairness}$ HDI included zero). However, Fairness improved the model's predictive accuracy according to the LOO-CV – although not the model's fit to the observed data (Table 3.1). To clarify whether this predictive accuracy improvement was because Fairness captured meaningful interindividual differences in our sample rather than collinearity between parameters, we analyzed only those trials in which Fairness and Joint-cost minimization predicted different decisions (6682 "unambiguous" trials).

People made co-efficient choices in 76.6% (*SD* = 10.4) of these trials, almost exactly the same proportion as in the full sample. We re-estimated the Self and Other Disparity and Fairness only models on this dataset and found that the former predicted the participants' behavior better (Table 3.1, Models 3.2 & 4.2). The estimates support the hypothesis that people made decisions that aimed to minimize both Self and Other costs (Figure 3.5; see Figure 3.10 for individual estimates). Overall, we found no clear effect of Fairness on decision-making, and conclude that the Self and Other Disparities model provides the most accurate description of our findings.

Table 3.1: Parameter estimates from all fitted models, with measures of predictive accuracy and model fit (LOO-CV – Leave-one-out cross-validation, AUC - Area Under the Curve). Coefficients were rescaled to express the effect of the cost disparities in units of one on-screen cm. Pixel-based estimates for all parameters are reported in Table 3.3.

Object choice (A1) predictors	**Mode of posterior [95% HDI]**			**LOO-CV [SE]**	**AUC**
	$\mu_{\beta Self}$	$\mu_{\beta Other}$	$\mu_{\beta Fairness}$		
(1) Self Disparity	-0.156 [-0.313, -0.005]			13109.1 [97.3]	0.735
(2) Other Disparity		-0.058 [-0.210, 0.092]		14529.2 [82.0]	0.628
(3) Self Disparity and Other Disparity	-0.367 [-0.469, -0.273]	-0.291 [-0.390, -0.193]		9437.8 [124.6]	0.859
(4) Fairness			-0.003 [-0.089, 0.092]	16153.9 [42.2]	0.542
(5) Self Disparity, Other Disparity, and Fairness	-0.376 [-0.475, -0.256]	-0.282 [-0.397, -0.185]	-0.031 [-0.120, 0.083]	9407.8 [124.5]	0.851
1. (3.2)Self Disparity and Other Disparity (unambiguous trials)	-0.348 [-0.440, -0.246]	-0.267 [-0.378, -0.146]		5414.6 [93.5]	0.855
2. (3.3)Self Disparity and Other Disparity (Block 1)	-0.298 [-0.414, -0.193]	-0.110 [-0.216, -4.08^{e-05}]		610.6 [27.2]	0.825
3. (4.2)Fairness (unambiguous trials)			0.172 [0.038, 0.318	7850.9 [75.0]	0.748

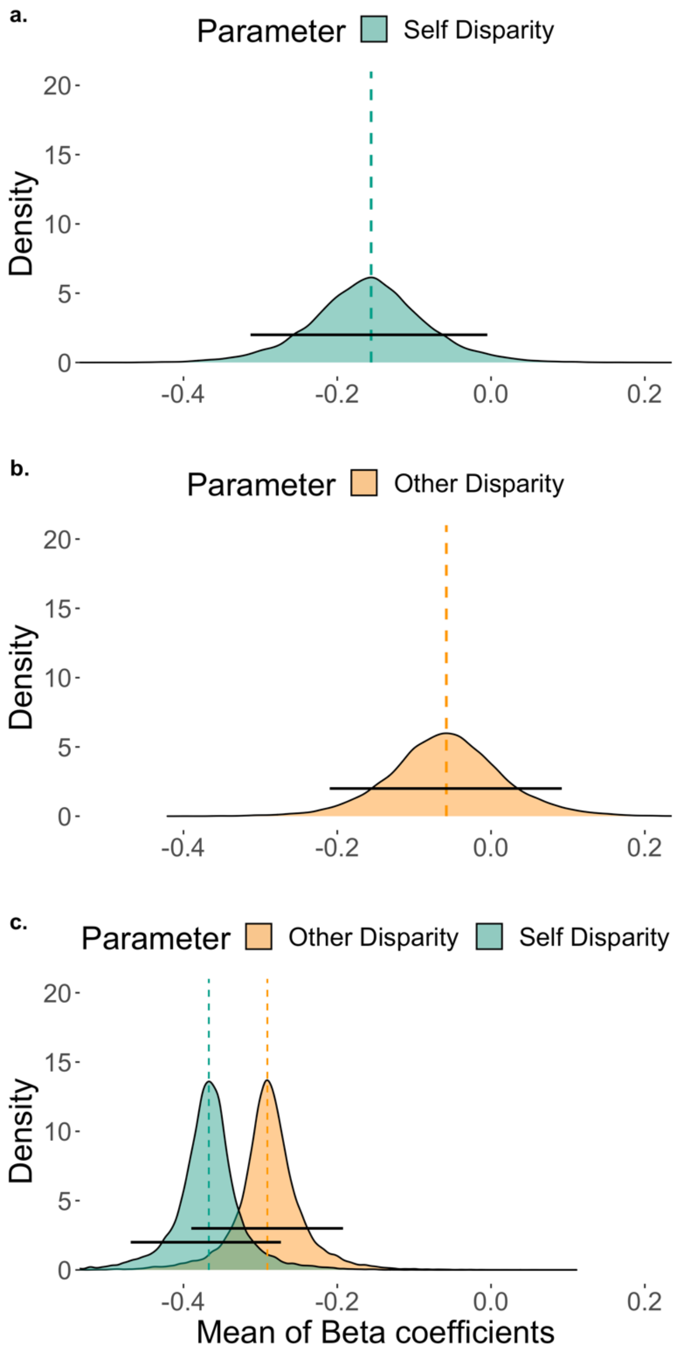

Figure 3.4: (a)-(c) Posterior probability distributions of the rescaled μ_β parameters for Self and Other Disparities in Models 1 to 3 (see Table 3.1), respectively. The dashed vertical lines indicate the modes of μ_β, the black horizontal lines represent the 95% HDIs.

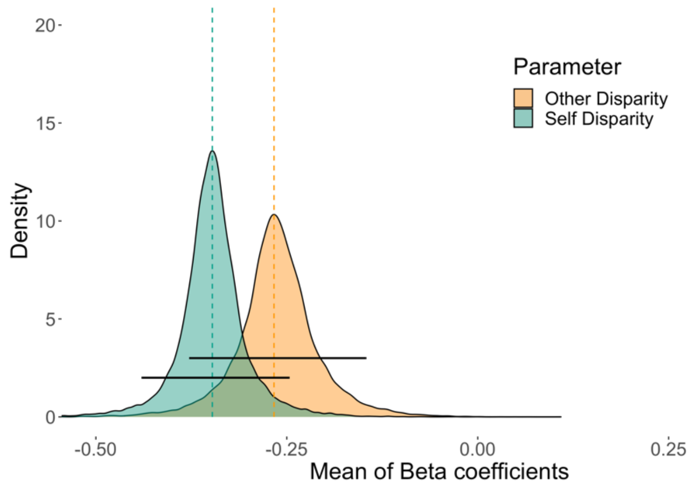

Figure 3.5: Posterior probability distributions of the rescaled μ_β parameters in the Self + Other Disparity model estimated on the non-ambiguous trials (Model 3.2). The dashed vertical lines indicate the modes μ_β, the black horizontal lines represent the 95% HDIs.

3.3.3 Learning and tit-for-tat strategy

Were participants' decision-making strategies stable over time? Did a partner's previous co-efficient choices drive behavior as a tit-for-tat strategy? We addressed these questions by running additional models, extended with the factors Trial and Block (of 5 trials), and a variable coding whether the co-actor chose co-efficiently on their previous trial (PrevCoeff). We found that neither predictor improved the best model's predictive accuracy or fit to the data (Table 3.5).

Additionally, we re-estimated all models on each participant's first 5 decisions (Block 1) and found that in the first minute of a game, decisions were best described by the Self and Other Disparity model, although with a higher relative weight on Self costs (Self: .73, 95% HDI = [.47, 1.00], Other: .27, 95% HDI = [0, 53]; Table 3.5, Figure 3.11). These results together suggest that participants adopted a stable Joint-cost minimizing strategy following a brief phase of relative Self-cost minimization (Figure 3.6).

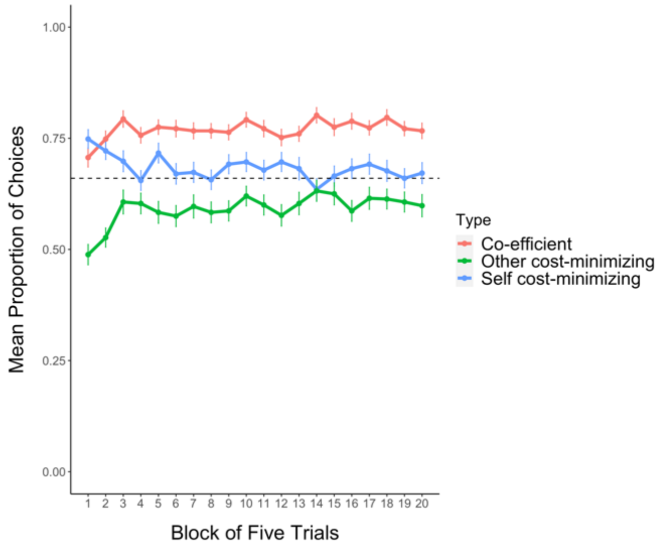

Figure 3.6: Mean proportions of Joint-, Self- and Other-cost minimizing object choices in 5-trial blocks. Each choice could be categorized in either or all of these categories due to overlaps between the predictions of each strategy in a trial. Error bars represent SE, the dashed horizontal line shows chance level.

3.3.4 Benefits of Co-efficient Decisions

To investigate whether co-efficient choices conferred any benefit to the dyad, we averaged the total trial durations for each dyad. When making co-efficient choices, dyads completed trials on average in M_{Means} = 4.55 s (SD_{Means} = 0.63 s), with an average standard deviation of M_{SD} = 0.98 s (SD_{SD} = 0.68 s). Trials with sub-efficient object choices numerically lasted longer for 55 out of 60 dyads, on average for M_{Means} = 4.78 s (SD_{Means} = 0.65 s), with a lower average standard deviation than in the case of co-efficient choices (M_{SD} = 0.83 s; SD_{SD} = 0.25 s). These results suggest a beneficial effect of co-efficient choices on task completion time. Statistical testing was not conducted due to the low number of sub-efficient choices.

3.4 Discussion

The current study explored the computations that underlie joint-cost minimizing decisions in planning joint actions. We tested the hypothesis that co-actors represent the collaborating dyad's joint costs as a sum of the members' individual costs and seek to minimize this value.

Participants made binary decisions between action plans with different associated movement costs in a joint object matching task. We modelled the parametric dependence of participants' choices on the action cost disparities for the acting participant and those for her partner using hierarchical logistic regressions. The cost disparities of the actor and co-actor were found to be almost equally weighted, which strongly suggests that the decision-maker aimed to minimize the joint action cost (77% of choices were co-efficient).

We also tested an alternative hypothesis, according to which the minimization of unfairness in the distribution of individual action costs determines action decisions. Overall, we conclude that fairness did not influence action choices. Furthermore, we did not find conclusive evidence for participants following a tit-for-tat strategy along the lines of "If you choose co-efficiently, I will do so, too". The results suggest that after briefly overweighting Self costs in the first few trials, the participants followed a Joint-cost minimizing strategy throughout the task.

Therefore, our experiments' results provide support for co-efficiency maximization as a primary strategy of action planning in a joint task involving two contributions. This is consistent with the way Kleiman-Weiner and colleagues (2016) operationalized cooperative planning in their computational model, 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 weighted sum of individual costs in joint action planning. Our findings suggest weights of ~.56 on the decision-maker's own, and ~.44 on her partner's individual costs. The behavior we found is also qualitatively consistent with previous findings from a joint action task (Török et al., 2019) and economic games (Colman et al., 2008a). Based on results from the co-representation literature (e.g., Schmitz et al., 2017), we speculate that in simultaneous tasks, too, people might account for joint costs.

It is possible that participants made choices based on hypothetical costs. However, we do not consider this a problem for our account. We argue that if participants had treated this as an abstract problem with more and less appropriate choices in a hypothetical mode (akin to a distance judgment task without movement), the fact that participants minimized movement distances for themselves and their partners suggests that the manipulation of even imaginary action costs influenced decision-making. This would strengthen our account, not weaken it.

Investigating the factors that might modulate how actors' individual action costs are weighed in decision-making awaits future research. Relevant factors might include the explicit role distribution of co-actors, social hierarchies (Kleiman-Weiner et al., 2016), the relative competence of the co-actors at specific motor tasks, and uncertainty about the co-actors' cost functions. Increasing the uncertainty about the partner's action costs might make people downplay the importance of a co-actor's individual costs in the computation of joint costs, or to ignore them altogether. Furthermore, more extreme costs or larger asymmetries between individuals might have similar effects on decision-making: the former could push people towards self-interest (and so could social hierarchy), the latter could inspire a fairness-focused strategy instead. The effects of these factors should be explored to 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 for a dyad in a sequential cooperative activity.

3.5 Supplementary Material

3.5.1 Methods – Additional Experiment-wise Information

3.5.1.1 Participant exclusions. We excluded dyads (1) due to computing errors caused by equipment failure, which on occasion resulted in multiple disruptions during data collection; or (2) when the correlational structure of the experiment's parameters was not as intended, due to the design's stochastic nature; or (3) when participants chose the same object in every trial. Table 3.2 shows the total number of participants per experiment and the specific reasons for exclusion.

Table 3.2: Excluded participants

Experiment	Total number of participants	Number of excluded participants	Reasons for exclusion
1	48 (24 dyads)	8 (4 dyads)	equipment failure (3); unintended correlation in parameters (1)
2	66 (33 dyads)	26 (13 dyads)	equipment failure (13) 7 dyads were included in the final dataset who whose sessions were disrupted once, but successfully resumed (results from this sample were consistent with those form the rest of the group)
3	52 (26 dyads)	12 (6 dyad)	equipment failure (1), lost data (4), participant compliance (1)

3.5.1.2 Design. The three experiments we conducted were different only in terms of the parameter pairs that were de-correlated from one another across trials. In Experiment 1, the individual Self and Other cost disparities were statistically independent from one another. In Experiment 2, Actor 1's individual costs were independent from the joint action costs; whereas in Experiment 3, the second actor's individual costs were independent from the joint costs. Here we describe the parameter sampling procedures applied in each experiment.

Experiment 1. In Experiment 1, our primary aim was to investigate the independent contribution of Self Disparity and Other Disparity to the actors' decisions. In order to make this possible, we kept the distributions of these two factors uncorrelated across trials. Thus, 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 pixels)[10]. We then randomly selected the positions of all objects in such a way as to match these disparities and with the constraints that (1) the distance between the starting positions and the objects be between 120 and 385 pixels[11], (2) the angle between the line from an object to the start position and the edge of the screen be a minimum of 15°, (3) the angular separation between the paths from the starting positions to the objects be at least 45°, (4) and the absolute distance between the objects on both sides be at least 124 pixels.

The sampling process that generated object arrangements guaranteed that Self Disparity and Other Disparity were uncorrelated (Figure 3.7a). However, as a direct consequence, Joint Disparity (the sum of the two individual disparities) had a triangular distribution and was positively correlated with both terms (Figure 3.7b-c).

[10] The script (exp1_layoutGen.m) is available on the OSF site of the project, along with the scripts for the other experiments: https://osf.io/r6mz3/?view_only=3f5fc782dac242adbe02bf3bc48158b0

[11] On the screen, 100 pixels were equal to approximately 5.3 cm.

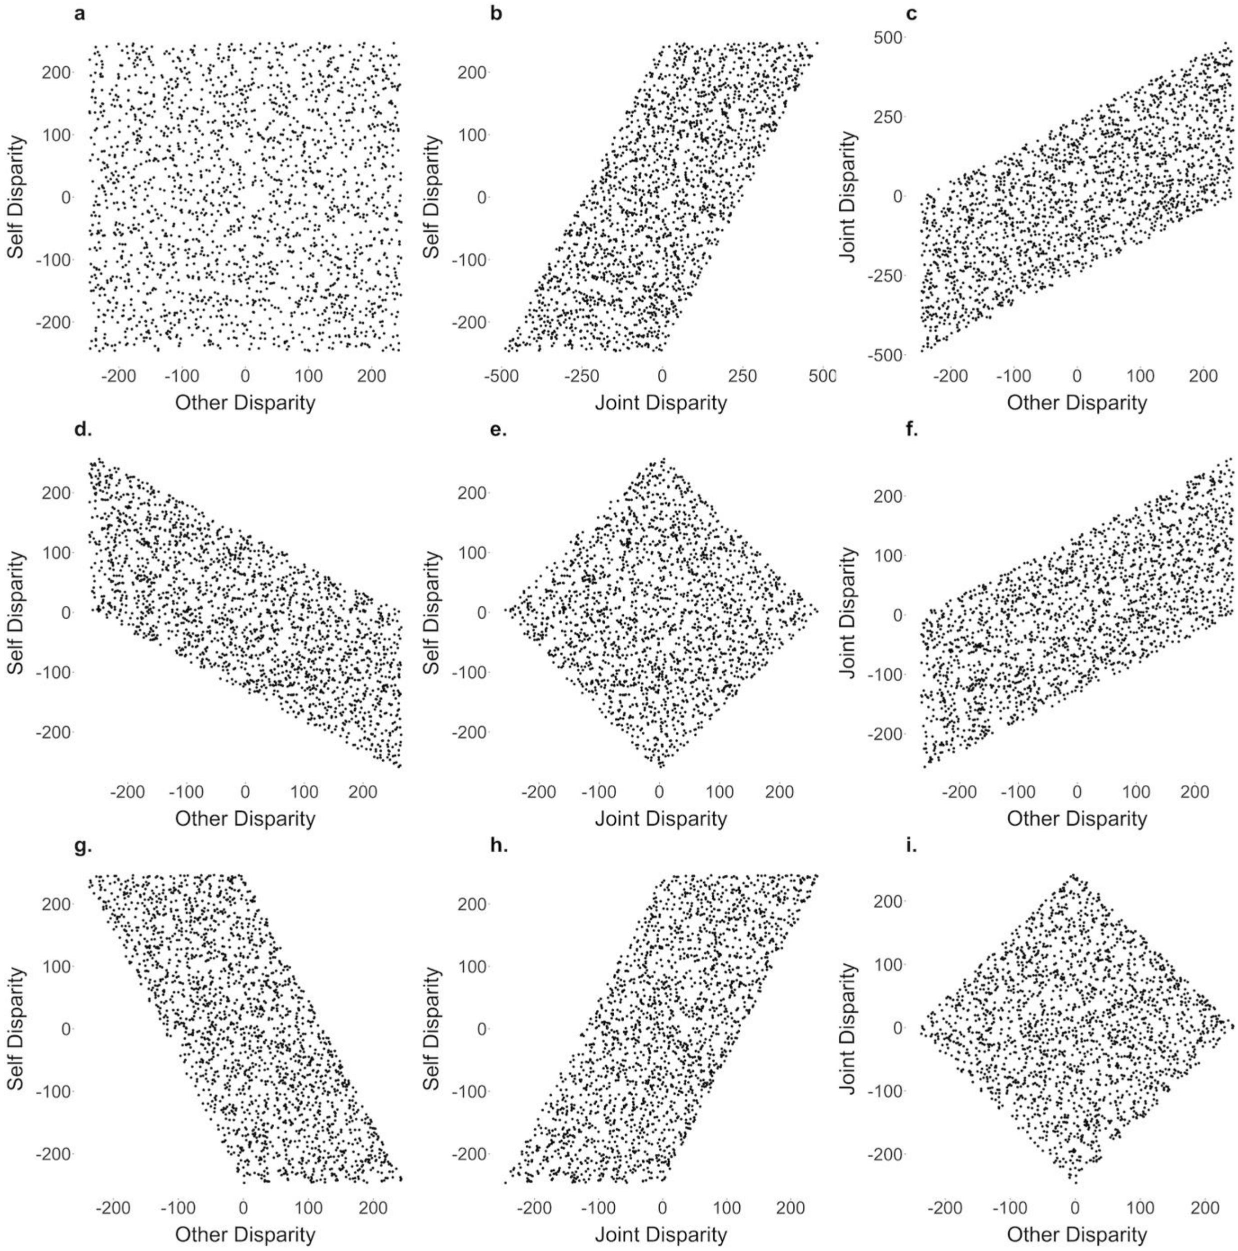

Figure 3.7: Scatterplots of the joint distributions of cost disparities in (a-c) Experiment 1, collapsed across all trials of all dyads (20 dyads, 2000 trials). (a) Self and Other Disparities were uncorrelated. (b) Self and Joint Disparities, and (c) Joint and Other Disparities were positively correlated with each other. (d-f) In Experiment 2, (d) Self and Other Disparities were negatively correlated with each other, (e) Self and Joint Disparities were uncorrelated, and (f) Joint and Other Disparities were positively correlated with each other. (g-i) In Experiment 3, (g) Self and Other Disparities were negatively correlated, (h) Self and Joint Disparities were positively correlated with each other, and (i) Joint and Other Disparities were uncorrelated.

Experiments 2 and 3. Experiment 2 tested the hypothesis that action initiators (Actor 1) plan their movements to minimize the summed aggregate action costs of the dyad's action sequence (*Joint Cost Disparity*) rather than to minimize their own individual costs (*Self Cost Disparity*). Experiment 3 probed the effect of Joint Cost Disparity against the individual costs of Actor 2 (*Other Cost Disparity*).

In both additional experiments, we applied the task from Experiment 1, and generated the layouts with the target objects' locations in the same way as in Experiment 1, with some important changes. We first sampled the individual – Self in Experiment 2, and Other in Experiment 3 – Cost Disparities for each trial from a triangular distribution with mode = 0 and limits provided by the maximum possible distance between an Actor's starting position and any target object (-265, 265 pixels). Then the parameters for Actor 2 (Other Disparity, Experiment 2) and Actor 1 (Self Disparity, Experiment 3), respectively, were drawn from a uniform distribution with limits set using the initially sampled Disparity parameter multiplied by -1.

Due to these sampling steps, the two actors' individual costs were negatively correlated with each other in both experiments (Figure 3.7d,g), and the Joint Disparity defined by the two individual parameters' sum was independent from the Self Disparity (and positively correlated with Other Disparity, Figure 3.7e-f) in Experiment 2, whereas it was independent from the Other Disparity in Experiment 3 (and positively correlated with Self Disparity, Figure 3.7h-i). As in Experiment 1, every dyad in both experiments completed 200 trials (100 uniquely generated trials per participant) in a pseudo-random order.

3.5.1.3 Description of the hierarchical model. We assumed that the trial-by-trial probability of choosing object A1 was Bernoulli distributed with parameter $\mu_{i|s,k}$, where *i* indexes the trial, *s* indexes the participant and *k* indexes the experiment that the participants participated in (see Trial level in Figure 3.8). The value of this parameter depended on a logistic function of the focal cost parameter(s) of the model weighted by the participant's β coefficient/s, $\beta_{Self,s,k}$, $\beta_{Other,s,k}$ or $\beta_{Fairness,s,k}$ (Subject level). The intercept was not estimated in the models, which is equivalent to assuming random decisions in the absence of any action cost disparities.

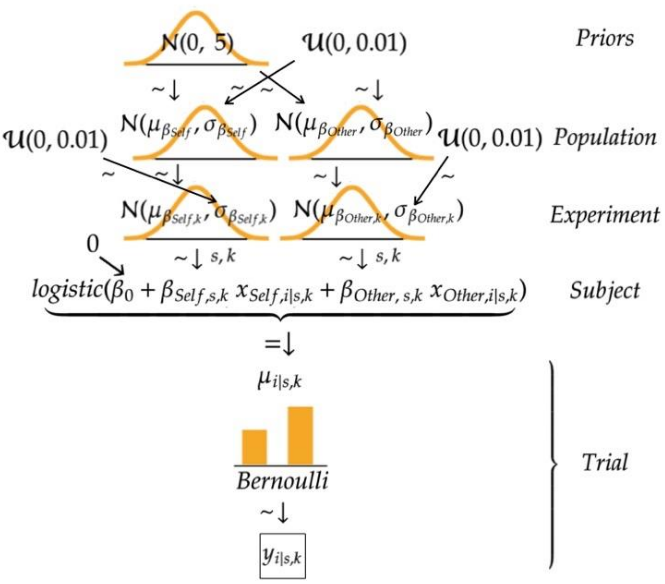

Figure 3.8: A graphical schema of the hierarchical regression model, adapted from Kruschke (2015).

The individual β coefficients were assumed to be normally distributed at the Experiment level around means $\mu_{\beta Self,k}$, $\mu_{\beta Other,k}$, and $\mu_{\beta Fairness,k}$ with standard deviations $\sigma_{Other,k}$, $\sigma_{Self,k}$, and $\sigma_{Other,k}$, corresponding to the assumption that participants' individual weighing strategies are noisy versions of a shared group-level weighing pattern within an experiment. We included a Population level above the Experiment level with μ_β and σ_β values for each cost parameter's β coefficients. Each experiment's $\mu_{Self,k}$, $\mu_{Other,k}$, $\mu_{Fairness,k}$, $\sigma_{Self,k}$, $\sigma_{Other,k}$, and $\sigma_{Fairness,k}$ parameters were assumed to be sampled from the Population level, e.g. $\mu_{\beta Self,k} \sim \mathcal{N}(\mu_{\beta Self}, \sigma_{\beta Self})$ and $\sigma_{\beta Self,k} \sim \mathcal{U}(0.0, 0.01)$. The priors for the Population level were set by hyperparameters $\mu_{\beta Self} \sim \mathcal{N}(0, 5)$ and $\sigma_{\beta Self} \sim \mathcal{U}(0.0, 0.01)$ (similarly for the other disparity parameters), a distribution around a zero effect of cost disparity. The priors for the σ_β parameters (and for the $\sigma_{\beta,k}$, parameters one level below) were set to approximately match the ranges of posterior σ_β estimates of the initial experiment-wise analyses[12] (Priors level). The same hyperpriors were used for all the predictors across all models.

12 N.B. Where comparison was possible, the experiment-level estimates did not qualitatively differ between the pooled analyses reported in the main text and in section 3.5.2 Results – Additional Experiment-wise Information (Table 3.3), and the original, experiment-wise, analyses (reported in Table 3.4). The original hyperpriors used for each experiment were $\mu \sim (0, 2)$ and $\sigma \sim (0.0, 0.5)$. See section 3.5.3 Additional Experiment-wise Information: Separate analyses for details.

3.5.1.4 Technical information on the estimation process. We customized Bayesian data analysis scripts that are freely available online to accompany Kruschke (2015)[13]. Specifically, we adapted a multiple logistic regression model (Kruschke, 2015, p. 622) to predict a categorical dependent variable (object choice) in a hierarchical structure, which enabled the simultaneous estimation of individual, experiment-, and population-level β coefficient distributions.

All models were estimated using a Gibbs sampler in the runjags package (Denwood, 2016) in R (version 3.5.1). Three chains were initialized using fixed seeds of three random number generators for the reproducibility of results. At first, 1,000 adaptation steps and 10,000 burn-in steps were taken and discarded before reaching convergence between the three chains. We kept 29,000 subsequent iterations for analysis, by thinning out every second step. Chain convergence was checked using Gelman and Rubin's (1992) convergence diagnostic, the potential scale reduction factor (PSRF). In most of the models, this factor's value was close to 1, i.e., chain convergence was satisfactory, and the full range of posterior distributions were explored. Although increasing the chain size would have ensured that all models' PSRF values be around 1, we had to compromise by capping the chain length at 29,000 iterations due to finite computational resources (to enable the calculation of WAIC and LOO-CV measures for model comparison, we had to estimate the log-likelihood at each trial, which placed considerable strain on our technical resources).

The data collected in the three experiments of the present study and the analysis scripts are available on the OSF site of the project:

https://osf.io/r6mz3/?view_only=3f5fc782dac242adbe02bf3bc48158b0

[13] The software and scripts were downloaded from https://sites.google.com/site/doingbayesiandataanalysis/software-installation .

3.5.2 Results – Additional Experiment-wise Information

We report the experiment-level parameter estimates for the eight logistic regression models reported in the main text. First, Table 3.3 summarizes these, together with the population-level estimates and measures of model fit; then follows a detailed description of the results of the five main models.

Table 3.3: Raw (pixel-based) parameter estimates and measures of predictive accuracy and model fit (WAIC – Watanabe-Akaike Information Criterion, LOO-CV – Leave-one-out cross-validation, AUC - Area Under the Curve) of the logistic regression models. Each row reports either population- or experiment-level estimates (indicated in the first column) for a given model.

Model	μ_β Mode	μ_β 95% HDI	σ_β Mode	σ_β 95% HDI	WAIC [SE]	LOO-CV [SE]	AUC
Model 1: Self Disparity					13108.7 [97.3]	13109.1 [97.3]	0.735
Population-level	-0.008	-0.017, -0.0003	0.006	0.003, 0.010			
Experiment 1	-0.012	-0.015, -0.010	0.009	0.007, 0.010			
Experiment 2	-0.002	-0.004, 0.0003	0.006	0.005, 0.009			
Experiment 3	-0.011	-0.013, -0.009	0.006	0.005, 0.008			
Model 2: Other Disparity					14529.0 [82.0]	14529.2 [82.0]	0.628
Population-level	-0.003	-0.011, 0.005	0.006	0.003, 0.010			
Experiment 1	-0.007	-0.010, -0.005	0.007	0.005, 0.009			
Experiment 2	-0.006	-0.007, -0.004	0.005	0.004, 0.007			
Experiment 3	0.004	0.001, 0.006	0.006	0.005, 0.008			
Model 3: Self + Other Disparity					9436.5 [124.5]	9437.8 [124.6]	0.859
Population-level	Self: -0.019	-0.025, -0.014	0.001	4.14e-07, 0.008			

	Other: -0.015	-0.021, -0.010	0.001	2.88^{e-07}, 0.008			
Experiment 1	Self: -0.019	-0.021, -0.016	0.009	0.007, 0.010			
	Other: -0.014	-0.017, -0.011	0.009	0.007, 0.010			
Experiment 2	Self: -0.019	-0.021, -0.016	0.006	0.005, 0.009			
	Other: -0.016	-0.019, -0.014	0.007	0.005, 0.009			
Experiment 3	Self: -0.021	-0.024, -0.018	0.007	0.006, 0.010			
	Other: -0.015	-0.018, -0.013	0.008	0.007, 0.010			
Model 4: Fairness					16153.8 [42.2]	16153.9 [42.2]	0.542
Population-level	-0.0001	-0.005, 0.005	0.002	0.001, 0.008			
Experiment 1	-0.001	-0.002, -0.0003	0.0001	1.03^{e-06}, 0.001			
Experiment 2	0.002	0.001, 0.003	0.003	0.002, 0.004			
Experiment 3	-0.001	-0.001, 0.0004	0.001	3.14^{e-07}, 0.002			
Model 5: Self + Other Disparity + Fairness					9406.3 [124.4]	9407.8 [124.5]	0.851
Population-level	Self: -0.020	-0.025, -0.014	0.001	3.81^{e-06}, 0.008			
	Other: -0.015	-0.021, -0.010	0.001	8.24^{e-08}, 0.008			
	Fairness: -0.002	-0.006, 0.004	0.001	8.98^{e-07}, 0.008			
Experiment 1	Self: -0.019	-0.022, -0.016	0.008	0.007, 0.010			
	Other: -0.014	-0.017, -0.011	0.009	0.008, 0.010			
	Fairness: -0.002	-0.004, -0.001	0.002	3.28^{e-06}, 0.004			
Experiment 2	Self: -0.020	-0.024, -0.013	0.006	0.004, 0.009			

	Other: -0.015	-0.022, -0.011	0.007	0.005, 0.009			
	Fairness: 0.001	-0.004, 0.005	0.0001	1.18^{e-07}, 0.003			
Experiment 3	Self: -0.021	-0.024, -0.019	0.007	0.006, 0.010			
	Other: -0.016	-0.018, -0.013	0.009	0.007, 0.010			
	Fairness: -0.002	-0.003, -0.0003	0.003	0.001, 0.005			
Model 3.2: Self + Other Disparity (unambiguous trials)					5412.4 [93.4]	5414.6 [93.5]	0.855
Population-level	Self: -0.018	-0.023, -0.013	0.001	3.99^{e-06}, 0.008			
	Other: -0.014	-0.020, -0.008	0.002	4.26^{e-07}, 0.009			
Experiment 1	Self: -0.017	-0.020, -0.014	0.008	0.006, 0.010			
	Other: -0.012	-0.015, -0.009	0.009	0.007, 0.010			
Experiment 2	Self: -0.019	-0.021, -0.016	0.006	0.004, 0.009			
	Other: -0.016	-0.019, -0.014	0.006	0.005, 0.008			
Experiment 3	Self: -0.019	-0.022, -0.017	0.006	0.004, 0.009			
	Other: -0.014	-0.017, -0.011	0.008	0.006, 0.010			
Model 3.3: Self + Other Disparity (Block 1)					606.9 [27.0]	610.6 [27.2]	0.825
Population-level	Self: -0.016	-0.022, -0.010	0.001	5.05^{e-07}, 0.008			
	Other: -0.006	-0.011, -2.16^{e-06}	0.001	3.26^{e-07}, 0.008			
Experiment 1	Self: -0.016	-0.021, -0.013	0.0003	9.35^{e-06}, 0.008			
	Other: -0.005	-0.009, -0.002	0.009	0.003, 0.010			

Experiment 2	Self: -0.015	-0.020, -0.010	0.009	0.002, 0.010			
	Other: -0.007	-0.010, -0.004	0.003	0.0001, 0.007			
Experiment 3	Self: -0.016	-0.020, -0.012	0.008	0.002, 0.010			
	Other: -0.006	-0.009, -0.0003	0.003	4.64^{e-05}, 0.009			
Model 4.2: Fairness (unambiguous trials)					7850.3 [74.9]	7850.9 [75.0]	0.748
Population-level	0.009	0.002, 0.017	0.004	0.003, 0.010			
Experiment 1	0.013	0.011, 0.015	0.003	0.001, 0.006			
Experiment 2	0.004	0.003, 0.006	0.004	0.003, 0.005			
Experiment 3	0.011	0.009, 0.013	0.004	0.001, 0.006			

3.5.2.1 Models 1 and 2: Self Disparity, Other Disparity. In Experiments 1 and 3, Actor 1's individual cost disparities exerted non-zero effects on the probability of their choosing object A1, whereas in Experiment 2, this effect was not different from zero. In Experiments 1 and 3, the modes of the posterior distributions of $\mu_{\beta Self,1}$ and $\mu_{\beta Self,3}$, the parameters for the experiment-level coefficient for the cost disparity, were -0.235 (95% HDI: [-0.291, -0.186]), and -0.204 (95% HDI: [-0.244, -0.164]), respectively. For every onscreen centimeter increase in Self Disparity, a 20.9% (Experiment 1) and a 18.4% (Experiment 3) decrease in the odds of an object A1 choice was expected. In Experiment 2, the estimated 95% HDI of the posterior distribution of the $\mu_{\beta Self,2}$ parameter included zero (Mode $\mu_{\beta Self,2}$ = -0.038, 95% HDI: [-0.078, 0.006]). The modal decrease in the odds of picking object A1 over B1 with a one cm increase in Self Disparity was 3.7%.

Model 2: Other Disparity. The estimation of the experiment-wise $\mu_{\beta Other,k}$ parameter's posteriors for the Other Disparity model found that Actor 2's cost disparity had negative, non-zero effects on the odds of object A1 choices in Experiments 1 and 2, whereas the estimates in Experiment 3 were distributed above zero.

In Experiments 1 and 2, the modes of the $\mu_{\beta Other,1}$ and $\mu_{\beta Other,2}$ parameter estimates were -0.137 (Experiment 1, 95% HDI: [-0.179, -0.097]) and -0.105 (Experiment 2, 95% HDI: [-0.137, -0.071]). This indicated that the expected decreases in the odds of an object A1 choice over a B1 choice, when Other Disparity increased by one centimeter, were 12.8% (Experiment 1) and 10.0% (Experiment 2). In Experiment 3, however, the posterior distribution of the Other Disparity's estimated $\mu_{\beta Other,3}$ parameter was fully above zero: with Mode $\mu_{\beta Other,3}$ = 0.066 (95% HDI: [0.026, 0.105]). This means that contrary to our predictions, increasing Other Disparity by one cm resulted in an expected increase of 6.8% in the odds of the participant picking A1 over B1.

To summarize, we found that in the case of the two single-predictor models, in 2 out of 3 experiments – when each of them was correlated with Joint Disparity – the disparities influenced decisions in the expected negative direction. The results of the estimations suggest that when each cost disparity parameter was de-correlated from the Joint Disparity of action sequences – i.e., Self Disparity in Experiment 2, and Other Disparity in Experiment 3 –, their effects were not as expected. Self Disparity by itself did not have an effect on choices (the 95% HDI included zero), whereas Other Disparity had an effect in the opposite direction than expected: when Other Disparity increased, the odds of an A1 choice also increased. This could possibly be due to an effect of self-cost minimization, because Other Disparity was negatively correlated with Self Disparity in Experiment 3.

3.5.2.2 Model 3: Self and Other Disparities. In all three experiments, the experiment-level means ($\mu_{\beta Self,k}$ and $\mu_{\beta Other,k}$) of the β_{Self} and β_{Other} coefficients for both disparities in Model 3 were distributed below zero. In Experiment 1, the mode of the $\mu_{\beta Self,1}$ parameter's posterior distribution was -0.355 (95% HDI: [-0.405, -0.306]), and the mode of the $\mu_{\beta Other,1}$ posterior was -0.265 (95% HDI: [-0.316, -0.207]). In Experiment 2, the two modes were similar in magnitude (Mode $\mu_{\beta Self,2}$ = -0.354, 95% HDI: [-0.394, -0.307]; Mode $\mu_{\beta Other,2}$ = -0.307, 95% HDI: [-0.356, -0.268]), as were the estimates in Experiment 3, although to a lesser degree (Mode $\mu_{\beta Self,3}$ = -0.389, 95% HDI: [-0.445, -0.341]; Mode $\mu_{\beta Other,3}$ = -0.290, 95% HDI: [-0.343, -0.244]).

Increasing Self and Other disparities (reported in this order) by a centimeter was expected to lead to a 29.9% and 23.3% decrease in Experiment 1, a 29.8% and 26.4% decrease in Experiment 2, and a 32.2% and 25.2% decrease in Experiment 3 in the odds of picking object A1 over B1. The 95% HDIs of the coefficients of the two cost disparities overlapped with one another in all three experiments, suggesting that there were no substantial differences between the magnitudes of the effects of the Self and Other disparities on decision-making. The relative average weights on Self and Other Disparity in the joint utility according to this combination model were .57 (95% HDI: [.49, .65]) and .43 (95% HDI: [.33, .51]), respectively, in Experiment 1; .54 (Self Disparity, 95% HDI: [.46 .60]) and .46 (Other Disparity, 95% HDI: [.41, .54]) in Experiment 2; and .57 (95% HDI: [.50, .66]) and .43 (95% HDI: [.36, .51] in Experiment 3.

3.5.2.3 Models 4 & 5: "Minimizing unfairness". The experiment-level estimates for the Fairness only model differed between the three experiments. In Experiment 1, we found a small non-zero effect of Fairness in the predicted direction (Mode $\mu_{\beta Fairness,1}$ = -0.024, 95% HDI: [-0.035, -0.006]). This suggests that with a one cm increase in the asymmetry in cost distribution between the two co-actors, a 2.4% decrease in the odds of an object A1 choice over B1 was expected.

In Experiment 2, we found a small effect in the opposite direction: the 95% HDI of the posterior distribution of the $\mu_{\beta Fairness,2}$ estimates did not include zero, with a mode of 0.032 (95% HDI: [0.011, 0.052]). This means that a one cm increase in cost distribution asymmetry related to object A1 resulted in a 3.2% odds increase of picking A1 object over B1. Finally, in Experiment 3, the 95% HDI of the posterior distribution of the $\mu_{\beta Fairness,3}$ estimates included zero ([-0.023, 0.008]), the most credible β coefficient was Mode $\mu_{\beta Fairness,3}$ = -0.009. These results suggest that overall, Fairness did not influence the probability of Actor 1 picking object A1 in a consistent manner across the experiments.

Model 5: Self, Other Disparity and Fairness. In the other combination model, we found similar patterns of results across experiments. In Experiment 1, the estimated posterior distributions of the Self and Other Disparity parameters' $\mu_{\beta Self,k}$ and $\mu_{\beta Other,k}$ values were both entirely below zero and the 95% HDIs of the two distributions overlapped with each other (Self Mode $\mu_{\beta Self,1}$ = -0.360, 95% HDI: [-0.411, -0.308]; Other Mode $\mu_{\beta Other,1}$ = -0.266, 95% HDI: [-0.316, -0.207]). In addition, we found a small non-zero effect of the Fairness parameter in the predicted negative direction (Mode $\mu_{\beta Fairness,1}$ = -0.042, 95% HDI: [-0.071, -0.018]). The expected odds decreases of an object A1 choice when each parameter was increased by one cm were 30.2% (Self Disparity), 23.3% (Other Disparity) and 4.1% (Fairness).

In Experiment 2, we found an even larger overlap between the effect sizes of Self and Other Disparity than in Experiment 1, suggested by a considerable overlap between the two 95% HDIs of the estimated posteriors (Self Mode $\mu_{\beta Self,2}$ = -0.371, 95% HDI: [-0.446, -0.254]; Other Mode $\mu_{\beta Other,2}$ = -0.288, 95% HDI: [-0.408, -0.216]). However, the effect of Fairness was not different from zero, suggested by the inclusion of 0 in the 95% HDI of the posterior distribution (Mode $\mu_{\beta Fairness,2}$ = -0.025, 95% HDI: [-0.081, 0.101]). By a one cm decrease in the Self and Other cost disparities, the odds of an A1 choice over the alternative decreased by 30.1% and 25.0%, respectively.

Finally, in Experiment 3, we found (similarly to Model 3) a small difference between the boundaries of the 95% HDIs of the posteriors of the μ_{β} coefficients on the Self and Other Disparities (Self Mode $\mu_{\beta Self,3}$ = -0.401, 95% HDI: [-0.454, -0.350]; Other Mode $\mu_{\beta Other,3}$ = -0.294, 95% HDI: [-0.347, -0.244]). This reflects the participants' tendency to minimize Self Disparities to a slightly larger extent than Other Disparities in Experiment 3 – which was also reflected in the positive coefficients on the Other Disparity in Model 2 (Other Disparity only model). Nevertheless, the results were similar to those in Experiment 1, in that the estimated beta weights on Self and Other Disparity were both much smaller than zero, in combination with a very small Fairness effect (Mode $\mu_{\beta Fairness,3}$ = -0.032, 95% HDI: [-0.057, -0.006]). Increasing each parameter by one cm resulted in expected decreases in the odds of an A1 choice by 33.0% (Self Disparity), 25.5% (Other Disparity), and 3.1% (Fairness).

Overall, the estimated weights on each parameter in the joint utility function, according to the combination model, were similar to one another across the three experiments. In Experiment 1, the weight on Self Disparity was .54 (95% HDI: [.46, .62]), on Other Disparity: .40 (95% HDI: [.31, .47]), and on Fairness, it was .06 (95% HDI: [.03, .11]). The values of these weights estimated based on Experiment 2's data were .59 (95% HDI: [.40, .70]) on Self Disparity, .45 (95% HDI: [.34, .64]) on Other Disparity, and .04 (95% HDI: [-.16, .13]) on Fairness. Finally, we found that in Experiment 3, the weights were .55 (Self Disparity, 95% HDI: [.48, .62]), .40 (Other Disparity, 95% HDI: [.34, .48]), and .04 (Fairness, 95% HDI: [.01, .08]).

3.5.3 Additional Experiment-wise Information: Separate analyses

We report in Table 3.4 the results of the original parameter estimations that we conducted on each experiment's data before pooling them together for the unified analyses. Since the designs of the three experiments differed in which pairs of cost disparities were de-correlated from one another, the models estimated also differed in the parameter combinations we used as predictors. Multiple-predictor models only included de-correlated parameter pairs; and in Experiments 2 and 3, we estimated only those single-predictor models of which the predictors were independent from Joint Disparity (i.e., in Experiment 2, we estimated only the Self Disparity model, in Experiment 3, only the Other Disparity model). In Experiment 1, although both individual cost disparity parameters were correlated with Joint Disparity, they were each tested as predictors in single-predictor models (Self Disparity only, Other Disparity only) to measure their predictive power against the combination model Self + Other Disparity.

3.5.3.1 Description of the experiment-wise hierarchical models. The experiment-wise models were identical in structure to the described pooled data model, except for the removal of the experiment level. We set the uninformed priors for this group-level distribution by vague hyperparameters ($\mu \sim (0, 2)$, $\sigma \sim (0.0, 0.5)$), 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 (Priors level).

Table 3.4: Raw (pixel-based) parameter estimates and model fit measures (DIC – Deviance Information Criterion, AUC – Area Under the Curve) of the original experiment-wise logistic regression models. The best fitting models' estimates for each experiment are set in bold.

Model	μ_β Mode	μ_β 95% HDI	σ_β Mode	σ_β 95% HDI	DIC	AUC
Experiment 1						
1: Self Disparity	-0.013	-0.016, -0.010	0.009	0.007, 0.012	3810	0.817
2: Other Disparity	-0.008	-0.010, -0.005	0.006	0.005, 0.009	4575	0.723
3: Self + Other Disparities	**Self: -0.019**	**-0.022, -0.016**	**0.009**	**0.006, 0.012**	**2467**	**0.898**
	Other: -0.013	**-0.016, -0.010**	**0.009**	**0.007, 0.012**		
4: Fairness	-0.001	-0.002, -0.0004	0.0002	3.57e-08, 0.002	5539	0.526
5: Joint Disparity + Fairness	Joint: -0.013	-0.014, -0.011	0.004	0.003, 0.005	3205	0.893
	Fairness: -0.002	-0.004, -0.001	0.002	3.92e-07, 0.003		
Experiment 2						
1: Self Disparity	-0.002	-0.004, 0.001	0.006	0.005, 0.008	5231	0.542
2: Joint Disparity	-0.016	-0.018, -0.014	0.006	0.005, 0.009	4023	0.822
3: Joint + Self Disparities	**Joint: -0.017**	**-0.020, -0.014**	**0.007**	**0.005, 0.010**	**3637**	**0.824**
	Self: -0.002	**-0.004, 0.001**	**0.008**	**0.006, 0.011**		
Experiment 3						
1: Other Disparity	0.004	0.002, 0.006	0.006	0.005, 0.008	5202	0.593
2: Joint Disparity	-0.019	-0.022, -0.016	0.007	0.005, 0.010	3821	0.843
3: Joint + Other Disparities	**Joint: -0.022**	**-0.026, -0.019**	**0.009**	**0.006, 0.013**	**3317**	**0.857**
	Other: 0.006	**0.003, 0.009**	**0.009**	**0.007, 0.012**		

3.5.4 Correlations between Perspective-Taking, Empathy, Liking a Co-actor and Behavior Data

It is possible that general abilities of perspective-taking and empathic concern in social interactions may prove useful in the computation of collective action costs in cooperative contexts. We report the results of exploratory correlational analyses (conducted on the pooled data) of the potential relationships between how much participants prioritized joint-cost minimization and their perspective-taking abilities and degree of empathy towards other people, as well as how much they liked their co-actors.

Following the object matching task and before being debriefed about the experiment, the participants responded to a short custom questionnaire on their perceived purpose of the study and how much they liked their partner ("How much did you like your co-player?"). Ratings of Liking the partner were obtained using a 7-point Likert scale (1 – Not at all, 7 – Very much). Participants also completed the Perspective-Taking and Empathic Concern scales from the Davis Interpersonal Reactivity Index (Davis, 1980) as measures of perspective-taking and trait empathy. The maximum score on both scales was 28.

To operationalize the weight that participants placed on minimizing the joint costs of an action sequence, we used each participant's proportion of co-efficient choices out of the 100 trials they completed as the decision-making Actor 1 ("co-efficiency ratio"). The higher the value of this measure, the bigger the weight a participant placed on minimizing the joint costs of an action sequence.

The average co-efficiency ratio was M = .77 (Range: .55 - .91, SD = .07). We found no statistically significant correlation between this measure and the Liking scores (Mdn = 6, interquartile range, IQR = 1, Spearman's ϱ = .080, p = .387). Likewise, we found no relationship between the co-efficiency ratio and either Perspective-Taking (Mdn = 19, IQR = 5, ϱ = -.038, p = .681) or Empathic Concern (Mdn = 20, IQR = 7, ϱ = -.064, p = .486). These results suggest that in the present task, joint-cost minimizing behavior was unrelated to the participants' perspective-taking or empathic abilities and to how sympathetic they found their co-actor.

3.5.5 Figures showing participant-wise estimates for the best-fitting models

The figures below present the individual-level parameter estimates according to the best-fitting model. Figure 3.9 shows the results of the analysis on the entire dataset, Figure 3.10 on the data subset where the predictions of fairness and co-efficiency were dissociated, and Figure 3.11 on the first Block of 5 trials.

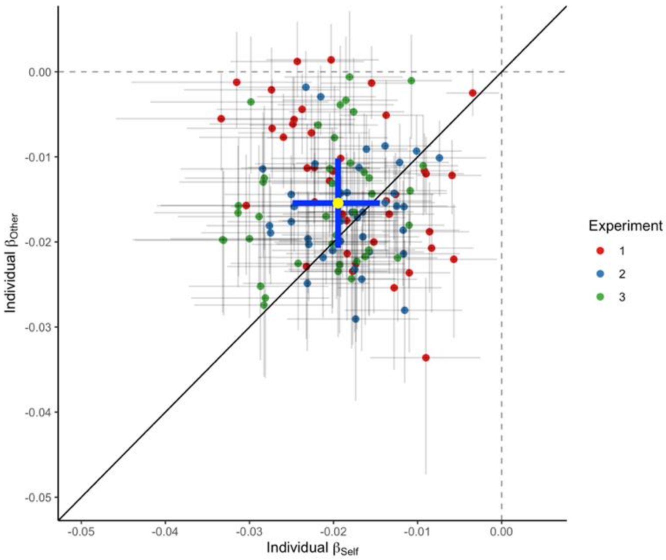

Figure 3.9: Individual posterior estimates of the raw $\mu_{\beta Self,k}$ and $\mu_{\beta Other,k}$ parameters in the Self and Other Disparity model (Model 3) run on the whole dataset of the three experiments. Each individual's posterior modes are shown with the 95% HDIs. The dashed horizontal and vertical lines indicate zero. The yellow point indicates the population-level modal $\mu_{\beta Self,k}$ and $\mu_{\beta Other,k}$ estimates, with the blue lines indicating their respective 95% HDIs.

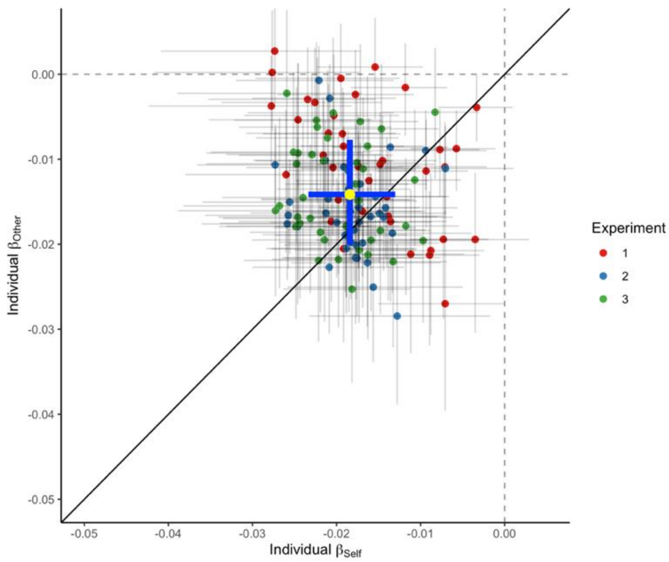

Figure 3.10: Individual posterior estimates of the raw $\mu_{\beta Self,k}$ and $\mu_{\beta Other,k}$ parameters in the Self and Other Disparity model (Model 3.2) run on the unambiguous trials only where the predictions of Fairness and Joint-cost minimization diverged. Each individual's posterior modes are shown with the 95% HDIs. The dashed horizontal and vertical lines indicate zero. The yellow point indicates the population-level modal $\mu_{\beta Self,k}$ and $\mu_{\beta Other,k}$ estimates, with the blue lines indicating their respective 95% HDIs.

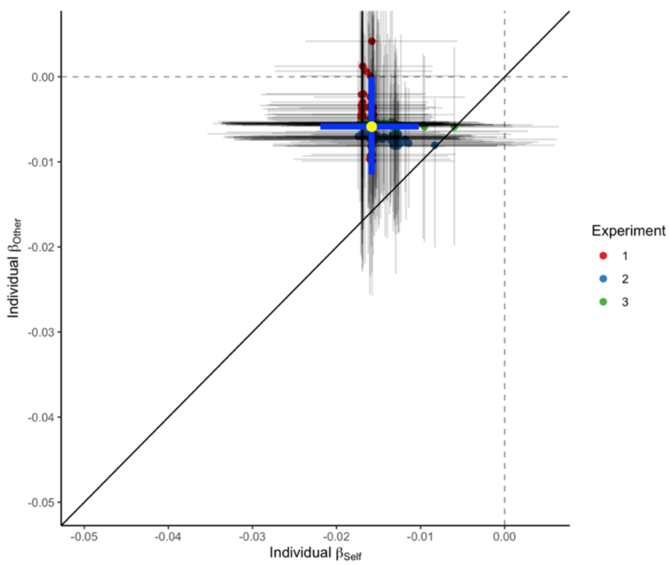

Figure 3.11: Individual posterior estimates of the raw $\mu_{\beta Self,k}$ and $\mu_{\beta Other,k}$ parameters in the Self and Other Disparity model (Model 3.3) run on the first 5 trials (Block 1) of each participant playing as Actor 1. Each individual's posterior modes are shown with the 95% HDIs. The dashed horizontal and vertical lines indicate zero. The yellow point indicates the population-level modal $\mu_{\beta Self,k}$ and $\mu_{\beta Other,k}$ estimates, with the blue lines indicating their respective 95% HDIs.

3.5.6 Additional models examining the effect of learning and reciprocity of co-efficient decisions on strategy use

Table 3.5: Measures of predictive accuracy and model fit (WAIC – Watanabe-Akaike Information Criterion, LOO-CV – Leave-one-out cross-validation, AUC - Area Under the Curve) of all of the logistic regression models mentioned in the main text. We include the 5 main models and the extended models addressing questions of learning and tit-for-tat decision-making.

Object choice (A1) predictors	**Whole Session**			**First Block (5 trials)**		
	WAIC [SE]	**LOO-CV [SE]**	**AUC**	**WAIC [SE]**	**LOO-CV [SE]**	**AUC**
(1) Self Cost	13108.7 [97.3]	13109.1 [97.3]	0.735	642.7 [24.8]	644.8 [25.1]	0.797
1. PrevCoeff + Self	12211.0 [109.8]	12211.9 [109.8]	0.776	636.2 [25.6]	640.3 [25.9]	0.803
Self * Trial	13711.5 [94.4]	13712.7 [94.4]	0.714	654.2 [26.1]	670.4 [27.9]	0.788
Self * Block	13653.1 [95.0]	13654.2 [95.1]	0.716	N/A	N/A	N/A
(2) Other Cost	14529.0 [82.0]	14529.2 [82.0]	0.628	801.6 [11.9]	803.8 [12.0]	0.513
2. PrevCoeff + Other	13416.9 [100.6]	13417.4 [100.6]	0.717	768.2 [16.0]	773.8 [16.2]	0.622
Other * Trial	14743.2 [80.8]	14744.0 [80.9]	0.627	797.5 [12.7]	804.0 [13.2]	0.513
Other * Block	14712.3 [81.2]	14713.0 [81.2]	0.627	N/A	N/A	N/A
(3) Self Cost + Other Cost	9436.5 [124.5]	9437.8 [124.6]	**0.859**	606.9 [27.0]	**610.6 [27.2]**	0.825
3. PrevCoeff + Self + Other	9443.3 [124.6]	9444.7 [124.6]	0.859	610.2 [27.2]	616.4 [27.5]	0.826
Self * Trial + Other * Trial	10220.0 [128.7]	10224.6 [128.9]	0.843	615.1 [28.8]	641.7 [31.2]	0.813

(Self + Other) * Trial	11864.1 [112.8]	11864.9 [112.9]	0.839	701.5 [21.5]	711.5 [22.4]	0.754
Self * Block + Other * Block	10119.0 [129.1]	10123.8 [129.4]	0.844	N/A	N/A	N/A
(Self + Other) * Block	11788.2 [113.6]	11789.0 [113.7]	0.840	N/A	N/A	N/A
(4) Fairness	16153.8 [42.2]	16153.9 [42.2]	0.542	830.8 [5.8]	832.6 [5.9]	0.521
4. PrevCoeff + Fairness	14627.3 [85.1]	14627.6 [85.1]	0.699	796.7 [13.0]	801.5 [13.2]	0.625
Fairness * Trial	16165.6 [42.9]	16165.9 [42.9]	0.543	831.5 [6.9]	836.1 [7.2]	0.521
Fairness * Block	16159.7 [43.0]	16159.9 [43.0]	0.544	N/A	N/A	N/A
(5) Self Cost + Other Cost + Fairness	**9406.3 [124.4]**	**9407.8 [124.5]**	0.851	609.0 [27.1]	614.7 [27.5]	0.818
5. PrevCoeff + Self + Other + Fairness	9413.3 [124.6]	9415.1 [124.7]	0.851	611.5 [27.3]	619.4 [27.8]	0.816
Self * Trial + Other * Trial + Fairness * Trial	10182.9 [128.9]	10189.2 [129.2]	0.836	606.1 [28.6]	645.3 [31.7]	0.804
(Self + Other + Fairness) * Trial	13763.6 [92.2]	13764.3 [92.3]	0.710	734.4 [19.5]	743.9 [20.4]	0.719
Self * Block + Other * Block + Fairness * Block	10080.4 [129.3]	10086.8 [129.6]	0.834	N/A	N/A	N/A
(Self + Other + Fairness) * Block	13721.2 [92.9]	13721.9 [92.9]	0.711	N/A	N/A	N/A
(6) PrevCoeff	15050.6 [76.7]	15050.7 [76.7]	0.703	798.7 [11.6]	800.9 [11.7]	0.619

Chapter 4. How do co-actors integrate the costs of different types of actions in joint action planning?

4.1 Introduction

Humans' ability to coordinate their actions with their conspecifics to reach shared goals has been the subject of a growing body of research in the last two decades (for reviews, see Knoblich et al., 2011; Vesper et al., 2017). A range of studies have looked at the ways in which people facilitate their co-actors' actions by making adjustments to their own movement kinematics (e.g., Dötsch & Schubö, 2015) or by selecting actions that reduce the effort that the co-actor would have to expend in completing their part of a sequential joint task (e.g., Gonzalez et al., 2011; Meyer et al., 2013; Ray & Welsh, 2011). Building on this line of research, we examined whether such facilitatory behaviors could be part of a broader joint action planning strategy aimed at reducing the collective efforts of acting (known as the shared-effort model, Santamaria & Rosenbaum, 2011; and as co-efficiency, Chapters 2-3).

In previous work, we observed shared effort-reducing, or co-efficiency maximizing, behavior in joint action sequences. In a binary choice task, people were instructed to transfer a football's image on a touchscreen between two diagonally opposite goal locations, by working together with a co-actor. They had to choose between two paths of different efficiency to complete the joint action. Participants selected individually inefficient actions if those contributed to an increase in the relative overall efficiency of the joint action (Study 1, Chapter 2). Furthermore, a follow-up study showed that actors computed expected joint action costs as the weighted sum of the individual costs related to each person's action options, with a slightly larger weight placed on the decision-maker's own cost than on the co-actor's (Study 2, Chapter 3). Taken together, these findings suggest that people make decisions in repeated interactions by taking the expected group effort into account, as an aggregate of individual efforts.

These conclusions are, so far, limited to situations where alternative action costs are directly comparable to each other on the same scale. In Study 2, every joint action sequence was a combination of two individual actions of the same type: Both co-actors had to continuously drag the image of an object to a target on a touchscreen so that cost was a direct function of distance. The present study aims to extend the scope of our investigation of co-efficiency to more naturalistic settings by addressing how people plan joint actions composed of different types of individual actions. Do people plan joint action sequences based on the judged relative costs of the individual contributions if this requires integrating costs of different action types? And if so, can they accurately determine which combination of individual contributions is least costly for the group?

As far as we can see, there are two necessary preconditions to this. Firstly, that people should be able (and motivated) to select actions that make them individually efficient, when the options are different action types. Secondly, a decision-making actor needs to take into account the difficulty of the task assigned to the other actor, which is supported by our own previous results and an extensive literature on co-representation and facilitation in joint action. Before spelling out the hypotheses of the present study, we briefly review the two strands of literature informing each precondition.

4.1.1 Individual action selection and judged relative costs

If we want to act efficiently when facing multiple different action options, we need to compare the expected costs of these options. However, even though comparing different types of actions such as riding a bike or walking somewhere is a ubiquitous issue, it is a non-trivial, 'apples-and-oranges' type of problem (Rosenbaum & Feghhi, 2019). Recent studies in experimental psychology have given insight into how people select certain actions from multiple alternatives when they act alone (Rosenbaum et al., 2013; Rosenbaum & Feghhi, 2019). Some of these studies investigated how actors decide between two sensorimotor actions like reaching for an object and walking while carrying it (e.g., Potts, Callahan-Flintoft, et al., 2018; Rosenbaum, 2008, 2012; Rosenbaum et al., 2011), whereas others addressed decisions between motor and cognitive tasks, like walking and counting or memorizing numbers (e.g., Feghhi & Rosenbaum, 2019, 2020; Potts, Pastel, et al., 2018; Rosenbaum & Bui, 2019).

These studies show that systematically exploring the choices that people make between two different action options using two-alternative forced choice (2AFC) tasks provides a simple but elegant way of estimating the subjective task difficulties, or judged relative costs (Rosenbaum & Gaydos, 2008), that underlie decisions to act. An important assumption shared between these studies is that people aim to minimize some measure of effort, a common currency between actions that enables the direct comparison of the alternatives. The term common currency was introduced by Rosenbaum et al. (2011) to action selection research from behavioral ecology (e.g., Cuthill & Houston, 1997), a field which aims to predict the macroscopic behavioral choices of animals. The common currency between different actions likely depends on multiple factors, for instance, the contexts of actions (Potts, Pastel, et al., 2018). For example, when comparing cognitive tasks in which errors might reasonably be made (e.g., memorization or counting tasks), the probability of success may become more relevant and might be used as common currency. On the other hand, some contexts emphasize time or burning calories, and therefore people may decide how to act based on these parameters (Potts, Pastel, et al., 2018). In yet other contexts of large-scale physical actions (e.g., walking), distance naturally seems to lend itself to movements in space as a relevant common currency.

Rosenbaum (2008) and Rosenbaum et al. (2011) set out to determine the relative costs of walking and reaching for an object and concluded that "a reasonable proxy for whatever the true common currency may be" (Rosenbaum et al., 2011, p. 136) between the two actions was *functional distance*. Functional distance was defined as an additive combination of walking and reaching distances, the latter weighted by a constant (Rosenbaum et al., 2011). In their experiments, participants were asked to choose between two possible combinations of walking to a table, reaching for a loaded bucket that rested on the table, and then carrying it to a goal area. The two action sequences could be executed along each side of a table standing perpendicular to the participant's frontal plane. The walking distances to the table from a start position, the reaching distances over the table to pick up the bucket, and the walking distances to the goal area were systematically manipulated to yield different overall combinations of walking and reaching distances. Participants were instructed to choose whichever action seemed easier to them.

The probabilities of choosing to walk along either side of the table were best predicted by psychophysical models based on functional distance (Rosenbaum, 2008; Rosenbaum et al., 2011). The fitted curves allowed the authors to calculate the relative cost of walking to reaching in the task: reaching was on average 10 times costlier than walking (Rosenbaum et al., 2011). Knowing this constant makes it possible to calculate the relative total costs (i.e., functional distances) of different action sequences. A follow-up study measuring the latencies of decisions between the same action alternatives suggested that cost comparison was likely based on parallel sampling of the action features, rather than on serial simulation of each potential course of action (Rosenbaum, 2012).

Sometimes, however, distance might not be the common basis for comparing actions. When choosing between a sensorimotor and a cognitive task, there are multiple contenders for a potential common currency for judging relative costs: the likelihood of errors (Feghhi & Rosenbaum, 2019; 2020), the likelihood of task sustainability (Rosenbaum & Bui, 2019), or time (Potts, Pastel, et al., 2018). In Feghhi and Rosenbaum's (2019) study, participants were required to choose between combinations of digit memorization and carrying a box through a gap. The lengths of digit lists and the widths of the gaps were manipulated, and the results suggested that participants tended to choose actions that reduced the likelihood of error (i.e., wider gaps and shorter lists, Feghhi & Rosenbaum, 2019). However, these results do not necessarily mean that error avoidance is the common determinant of task difficulty. In a follow-up study, Feghhi and Rosenbaum (2020) found that although participants chose actions that reduced the likelihood of errors in both the mental and physical domains, error avoidance by itself was not sufficient to explain the observed behaviors.

Potts, Pastel and colleagues (2018) tested the hypothesis that the common currency that people use to decide between the mental task of counting and the physical task of carrying a bucket was the time spent on each task. Their results suggested that subjective task durations explained decisions the best, which they argue aligns with the idea that subjective time is one proxy for task demands in metacognitive evaluations of task difficulty (Dunn et al., 2016; Potts, Pastel, et al., 2018). More recently, Rosenbaum and Bui (2019) tested the alternative hypothesis that the judged sustainability of tasks (i.e., a judgment of whether the actor could perform the given task X, Y, or Z times) provides a better ground for comparing action costs than time, but predictions of the sustainability model significantly differed from the observed behavior. Time seemed to be a likelier index of task difficulty (Rosenbaum & Bui, 2019).

The research described above from Rosenbaum and colleagues show that comparing the relative costs of two actions of different kinds is a non-trivial task, and there is not one common currency based on which the comparison may be made. Aside from the rich exploration of the possible factors that guide decision-making in individual action planning, these studies also offer an important methodological takeaway. Action planning research benefits greatly from the use of psychophysical methods and 2AFC tasks when trying to address everyday apples-and-oranges problems. The psychometric curves that predict binary choices as functions of a hypothesized common currency (e.g., functional distance) provide approximations of the relative costs of two potential actions in simple numerical terms.

4.1.2 Co-representation of task constraints and difficulty

Beside planning actions efficiently for oneself based on the judged relative costs of the available options, in cooperative contexts, an actor would need to consider the difficulties of the action options available to his or her partner, too. In social interactions, actors may represent multiple aspects of the task at hand, for example, the identity of the actor who should be acting at a given time in a task (Philipp & Prinz, 2010; Wenke et al., 2011), or the stimulus-response (S-R) mappings assigned to the co-actor, referred to as *task co-representation* (Sebanz et al., 2003, 2005).

Task co-representation aids the prediction of a co-actor's action even if her actions are unseen[14], which contributes to the appropriate planning of one's own actions (Vesper et al., 2013).

Schmitz et al. (2017) tested the hypothesis that actors represent their co-actors' specific task constraints in a temporal coordination task by measuring the kinematic features of the actor's movements under varying relative constraints. In a modified version of van der Wel and Fu's (2015) task, pairs of participants were required to move a dowel rod back and forth between two circular targets each, with the joint goal of synchronizing their landing times on the target locations. In one type of trials, one of the participants had to move her dowel over a cardboard obstacle (constrained actor), whereas the other participant had no obstacle in the way of his movement (unconstrained actor).

Confirming the authors' predictions, the unconstrained actor's peak movement height was larger in trials where the partner had to clear an obstacle than in No Obstacle trials. This height increase reflected individually inefficient performance relative to the unconstrained actor's individual movement baseline. Visual access to a partner's actions seemed not to influence this effect, suggesting that the participants represented their co-actor's task, and the modulation of their movements was not the result of purely perceptual mechanisms such as visuomotor interference (Brass et al., 2000; Schmitz et al., 2017). Schmitz and colleagues (2017) also ruled out the possibility that participants co-represented the kinematic parameters of the actions performed by the co-actor, rather than the constraints of clearing the obstacle. The findings suggested that the participants represented the height of their co-actor's obstacle.

Similarly, Vesper et al. (2013) found that specific knowledge of a partner's task – jumping distance - helped synchronization of an actor's movements with an unseen co-actor via predictive simulations. In their task, participants had to synchronize forward jumps with a partner on the other side of a partition. Although the co-actor's movement could not be observed, before starting

[14] But see Böckler et al. (2012), the motor condition in Elekes et al. (2016) and Welsh et al. (2007) for examples of co-representation effects disappearing when participants had no visual access to the co-actor's actions.

to move, both participants received information about their own and their co-actor's task in a given trial, i.e., how far they had to jump forward. The participants modulated their movement preparation and execution phases according to the differences between task difficulties across co-actors. The larger the difference between the instructed distances, the later and higher the participant with a relatively shorter jumping distance ("easier" task) started to jump, relative to the participant with the longer distance ("harder" task). On the other hand, the participant with the longer jumping distance in a trial generally sped up their movement preparation phase. Overall, the effort distribution in coordination was task-specific, showing that planning actions with a goal to coordinate benefitted from the co-representation of the relative task difficulties (Vesper et al., 2013).

What other kinds of task constraints might actors co-represent? In a series of experiments, Schmitz et al. (2018) showed that task co-representation can also contain the order of movements to be executed in a co-actor's action sequence. Participants were required to synchronize the final steps of their action sequences, in which (similar to Schmitz et al., 2017) they had to place a dowel rod at target locations in pre-specified orders. According to the predictions of the task co-representation account, when a co-actor's sequence specified a different order of the same sub-tasks as the participant's, movement times should reflect an interference between the different task representations. This is what Schmitz et al. (2018) found for different kinds of task constraints including the horizontal distances between the targets, and the index of difficulty of the movements over fixed distances (using Fitts' law to manipulate target sizes; Fitts, 1954).

To sum up, these studies suggest that in simultaneous joint actions with the shared goal of temporal coordination, actors benefit from representing their co-actors' exact task conditions (e.g., how high the hand has to move to clear an obstacle, Schmitz et al., 2017), and the relative task difficulties between the co-actors (e.g., how far each actor has to jump, Vesper et al., 2013). By using this information, participants can modulate their behavior to achieve coordination.

From the perspective of co-efficiency, sequential joint actions provide a more easily tractable testing ground than simultaneous actions (e.g., by preventing the online influence of a co-actor's unfolding action on the decision-maker). In the context of sequential tasks, studies that focus on the facilitation of a partner's actions provide evidence for the co-representation of a co-actor's expected task difficulty (Ray et al., 2017), end-state comfort[15] (Dötsch & Schubö, 2015; Meyer et al., 2013), and beginning-state comfort (Gonzalez et al., 2011). These studies focus on the behavior of individuals initiating an action that reduces the effort of another actor finishing the sequence, suggesting that the initiator must hold an approximate representation of the partner's effort, or at least be sensitive to cues of effort.

Facilitation is often scaled to the partner's task constraints and action capabilities. For example, participants in a joint pick-and-place task modulated their initial grip on and rotations of mugs as necessary, when their co-actors had to place the mug in a cued final position (Dötsch & Schubö, 2015). The more rotation the final position would have required from the partner, the more the initiator executed of this rotation, despite not being instructed to do so. The authors interpreted this as evidence that participants represented both their own and their co-actor's tasks and planned their own actions to ensure the partner's end-state comfort by facilitating their movements in advance (Dötsch & Schubö, 2015). Further evidence for joint action planning that accommodated a partner's comfort was found in studies based on the selection of hand and grasp types in an object manipulation task (Scharoun et al., 2016), the adjustments of a water jug's (Ray & Welsh, 2011) and a mug's handle orientation (Constable et al., 2016), choosing grasp locations on an object (Meyer et al., 2013), and choosing spatial locations for handing over objects (Gonzalez et al., 2011; Ray et al., 2017; Scharoun et al., 2017).

[15] The end-state comfort effect refers to people's tendency to select actions that conclude in comfortable final limb postures (Rosenbaum et al., 1996). The beginning-state comfort effect, on the other hand, refers to posture selection that ensures that an action starts in a comfortable limb configuration. Gonzalez et al. (2011) found that people plan their actions involving object manipulation in a way that ensures their own end-state, and their interaction partner's beginning-state comfort, if the partner is to use the object.

To the best of our knowledge, no studies have addressed so far whether and how an actor represents the difficulties of multiple potential actions of a co-actor when these involve different types of actions. However, Ray et al. (2017) investigated how people performing the first half of a joint action sequence represent different aspects of their interaction partner's action and task difficulty when there is uncertainty about the exact action that the partner will have to perform. The authors predicted that if an action initiator represents the difficulty of her co-actor's potential action, this would be reflected in response selection in the following ways: the first actor would place an object to be manipulated by the co-actor 1) in a location that equalizes the index of difficulty (Fitts, 1954) of the potential actions, and 2) further into the contralateral space of the co-actor, which would facilitate their subsequent movement (Ray et al., 2017).

In each trial of the sequential aiming task, the participant first had to move a wooden dowel from a home position to a location between two targets, after which the co-actor moved the dowel to one of the targets (Ray et al., 2017). The target sizes were manipulated between and within trials, resulting in different indices of action difficulty. Crucially, during the first actor's action planning, the target's identity was unknown, and was only signaled by a cue after the dowel was grasped by the second actor. The results showed that the participants tended to place the dowel closer to the smaller targets than the larger ones, that is, they moved in ways that equalized the index of difficulty of the second actor's two potential actions in a trial (Ray et al., 2017). This would minimize the partner's expected action costs. Additionally, participants had a relatively weaker tendency to place the dowel into the partner's contralateral space, making movements easier for them. Ray et al. (2017) interpreted the observed behavior as evidence for sufficiently detailed representations of a co-actor's future action difficulty based on multiple task features.

To summarize, previous research has shed light on various aspects of shared task representations in social interactions, be they task constraints in simultaneous joint actions where the goal is temporal coordination (Schmitz et al., 2017, 2018; Vesper et al., 2013), or constraints in sequential joint tasks where people modulate their facilitatory behaviors according to how effortful or uncomfortable their partner's action would be (e.g., Dötsch & Schubö, 2015). People also seem to represent the index of difficulty of their partner's potential actions when not knowing the exact constraints the partner will have: They modulated their action by taking into account the uncertainty of the partner's future effort (Ray et al., 2017). Overall, the literature reviewed above suggests that individuals decide between actions of different kinds by minimizing some hidden common currency (section 4.1.1 Individual action selection and judged relative costs), and that co-actors represent their partner's task constraints, actions and expected efforts (section 4.1.2 Co-representation of task constraints and difficulty).

4.1.3 Research question and hypotheses

It remains an open question how actors planning a joint action sequence might integrate the costs of different action types in co-efficient planning. When faced with multiple different action types, do individuals choose actions that minimize some common currency behind the two actions, e.g., the time spent performing an action or a variable related to time, such as functional distance or effort? If they do so, do they integrate their partner's estimated relative costs into the planning process to plan co-efficient action sequences?

First, we hypothesize that individual action choices will exhibit a consistent pattern as a function of some variable (the hidden common currency) that monotonically changes with the experimental manipulations, be it time or some correlate of it[16]. However, we do not want to make strong claims about what the common currency underlying action selection might be in our study, since its identity is not a focal question for the investigation regarding joint action planning. The focus of the study is ultimately on the exploration of decision-making in social contexts. Since the

[16] Based on Rosenbaum and Bui (2019) and Potts, Pastel, et al. (2018), we could hypothesize that using time as a proxy for the common currency between actions would provide an efficient way to solve planning problems in both individual and joint actions, as time is an amodal quantity (Potts, Pastel, et al., 2018). The passing of time may also be potentially easier to perceive than a co-actor's physical effort in completing relatively simple movements, as time itself does not depend on the co-actor's movement skills but effort does. Equally, we could also hypothesize that in sensorimotor joint tasks using different types of hand movements, as in the present task, functional distance (Rosenbaum et al., 2011) may play the role of common currency. We assume that functional distance is related to time, i.e., movement duration (the longer the distance, the longer the duration), and since we cannot a priori define functional distances for the individual task, we use movement duration as a proxy for the common currency when examining individual choices.

duration of an action is a readily available measurement, we will use this as a proxy for the common currency to test the first hypothesis regarding individual efficiency. We predict that in an individual setting, people will make decisions that minimize the time spent on an action when choosing between two different types of sensorimotor actions.

Secondly, we hypothesize that to the extent that actors make efficient (consistent, time-minimizing) choices between different types of actions in an individual setting, they will also take into account the potential costs of multiple action possibilities of their partner when planning joint actions. We expect that people will approximate a co-efficient strategy, based on their estimations of the potential individual relative action costs for both co-actors, operationalized as functional distances[17] in pixels.

4.1.4 The experiment

To test these hypotheses, we designed an experiment that first measures how individual actors decide between tapping and dragging actions on a touchscreen (Figure 4.1a provides an example of the layout used in the task), and then tests how they plan when the same tasks are embedded in a social context (Figure 4.1b). We chose tapping and dragging actions to ensure that the two options in each trial were different enough so that their relative costs could be more pronounced on certain trials, and less pronounced on others. Even though both movements classify as discrete, reaching a target relies on different overall actions in each case. That is, in the case of tapping (which is a ballistic action based on its short duration, Gan & Hoffmann, 1988), multiple taps were required to reach a target object, whereas only one continuous dragging movement (non-ballistic) was necessary for a dragging action. This means that while in the latter case the movement had only one overall acceleration-deceleration profile, due to repetition in the former case, multiple acceleration-deceleration phases occurred in a trial separated by short breaks

[17] In the present task, we fixed the physical distances between the two potential targets and the starting locations on the screen within a person in a trial – therefore raw onscreen distance as a visual cue for cost could not directly be used as a potential common currency between two actions. Hence the proposal of functional distance to test the joint hypothesis, following Rosenbaum et al.'s (2011) example.

(i.e., lifting the finger from the screen). In addition, tapping sequences were aimed at a specific area on the screen, whilst dragging actions were not constrained spatially by instruction[18].

On the other hand, the two actions are not so different from each other that they would inspire participants to only choose the same action on every trial. In a pilot experiment ($N = 11$) we confirmed that tapping and dragging actions posed similar challenges to the participants based on a lack of difference in the grand average times spent performing them. At the same time, these actions were sensitive to experimental manipulations. Based on the pilot results[19], we could reasonably expect that manipulating the lengths of distances to be covered on the screen (Path) and the number of taps required (Step Number) would result in trials where participants would follow a strategy of making trial-by-trial decisions between actions, rather than always choosing the same type of action over the other regardless of Path or Step manipulations. This is because under some combinations of the Path and Step factors in a trial, dragging was a faster action than tapping, whereas with other Path-Step combinations, tapping was the faster solution.

The experimental design follows the logic of the design of Study 2 (Chapter 3), where we examined how people aggregate expected action costs that are expressed on the same scale: distance. Logistic regression models were fitted to decision data and the models were compared to each other to determine whose action cost influenced decision-makers' choices in the action planning process – the decision-maker's, the co-actor's, or both actors'. The present study was designed to ultimately enable a similar comparison of the cost disparities of two different action types, following a transformation of the numbers of taps required into *functional distances in pixels*. Distance is related to the time spent on an action, therefore, we regard it to be an appropriate operationalization of cost for our purposes.

[18] Participants were explicitly told they were free to follow their own trajectory between the starting location and the target objects.

[19] In trials with only 3 steps to be taken over a given distance, participants' movement durations were shorter when tapping than when dragging over that path (regardless of its length); while in trials with more steps (5 and 6), dragging resulted in shorter movement times than tapping. Path length affected movement durations in the opposite way. Over shorter distances (200 px), dragging resulted in shorter movement times than tapping, and for longer distances, tapping seemed to yield slightly faster movements than dragging.

We used a within-subjects design. First, in an Individual No Choice condition, we measured participants' movement times – this allowed us to make predictions about minimizing time. Then, we tested these predictions in an Individual Choice condition, where we estimated individual baseline cost functions for each participant by fitting psychometric curves to their decision data, similar to the cost estimation strategy of Rosenbaum (2008) and other studies from his lab (e.g., Rosenbaum et al., 2011). We calculated the participants' thresholds for switching between the two action types by computing their points of subjective equality on the curve (PSE, Knoblauch & Maloney, 2012; Lee & Wagenmakers, 2013). In psychophysics, the PSE denotes the point along a stimulus dimension where a stimulus intensity (e.g., the loudness of a sound) is judged by a person to have equal intensity to a standard stimulus (Lee & Wagenmakers, 2013). In the present case, we define the PSE as the amount of on-screen displacement of the participant's octagon granted by one tapping action, where the participant perceived the two action options as equally costly. Finally, we used the estimated PSEs to transform all tapping costs to the pixel-scale, which provided the functional distances of taps.

4.2 Methods

4.2.1 Participants

The target sample size was set to N = 40. This was based on the sample size in Study 2, which found cost minimization effects using a similar paradigm, suggesting sufficient statistical power for 20 dyads. We recruited participants through Central European University's Research Participation System and a student job agency. People who took part in the study reported in Chapter 3 did not participate in the present experiment. We tested only participant pairs who did not know each other. The participants gave their informed consent and received vouchers in exchange for their participation. The study was approved by EPKEB, the Hungarian ethics committee for psychological research.

Fifty right-handed participants were tested in total, and the data of 40 people were analyzed (26 females, 11 males, 3 did not specify; age $M = 23.0$ years, $SD = 4.31$). We excluded the first dyad we tested because one of the two participants arrived considerably earlier than their co-actor and had to wait, which could have affected their cooperative behavior in the joint condition (see Sacheli et al., 2012, on the effects of a negative interpersonal relationship on coordination). Three dyads were excluded because of equipment failure, and one dyad was excluded because of experimenter error.

4.2.2 Apparatus

The task was performed on a touchscreen monitor (Iiyama PROLITE 46", resolution 1920 X 1080 pixels, separate sync - horizontal: 31.47 – 67.5 KHz, vertical: 47 – 63 Hz)[20] 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 (Brainard, 1997; Pelli, 1997; Kleiner et al., 2007) in MATLAB® (The MathWorks, Natick, MA). Two response boxes (Black Box Toolkit Ltd.) were used to control trial onset and to register object choices. The participants used three keys on the response box (which has four keys in total). The middle key was aligned with the mid-line of the screen's longer side, and the other two keys were located equidistantly from the middle key (± 4.75 cm from the center of the middle key). The distance between the edge of the screen and the center of the middle key was 7 cm.

4.2.3 Stimuli and Task

In the individual conditions, one black square and one black circle (30 x 30 px) were displayed on the touchscreen, both positioned in the screen area close to the participant (Figure 4.1a). In the joint condition the screen was divided in two task areas, and two pairs of black squares and black circles were distributed between co-actors (Figure 4.1b). A green octagon (64 x 64 px) was used as a "collector" tool to collect one of the black target objects. At the start of each trial, the octagon

[20] This was the same monitor that we used in Study 2.

was orange-colored and positioned in a larger, black-bordered octagonal starting location (96 x 96 px) located mid-way along the longer side of the screen, aligned with the response box's middle key. The orange octagon signaled which participant would initiate the joint action in the role of Actor 1. After 3 seconds, the color switched to green, to cue the participant to start moving.

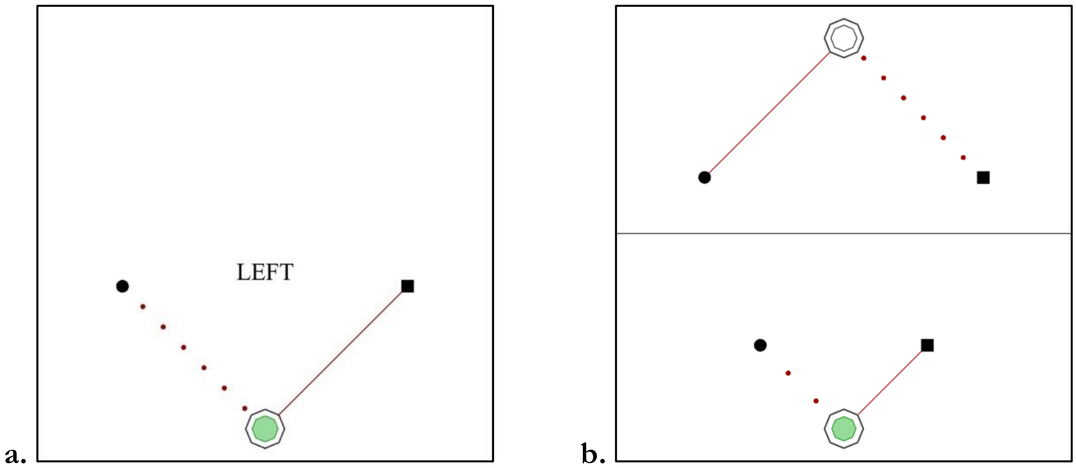

Figure 4.1: Examples of layouts in (a) one of the two individual conditions and (b) the joint condition. The image in (a) shows a trial in the Individual No Choice condition, where the participants were instructed to collect a specific target object referred to by its on-screen position (in the Individual Choice condition, the trials looked the same, without any instruction displayed). The dark red lines and dots indicate the type of action to be performed to reach each of the two black target objects, respectively (dots: tapping, line: dragging).

In the task, the participants had to execute one of two different kinds of actions to collect the black target objects: *tapping* repeatedly in the black-bordered octagonal starting location on the touchscreen to move the green octagon toward the target or *dragging* the green octagon onto the object with the index finger. The two action options were clearly indicated visually on the screen to the participants by different cues. In the joint task, the two options for each co-actor were shown in separate halves of the screen, with the octagons closest to each co-actor, and the tapping and dragging trajectories were always oriented diagonally from one another across the screen (Figure 4.1b). That is, from both co-actors' perspectives, the same type of action was assigned to the same side of space: for instance, in Figure 4.1b, tapping trajectories were located by the left, and dragging trajectories were located by the right hand side for both actors.

Tapping actions were signalled by the use of dark red dots (12 x 12 px). The number of dots in a trial between the starting location and a black target object indicated the number of taps required to reach the object, minus one (the one additional tap was executed for the final "jump" between the last red dot and the target). Going forward, we will refer to the total number of taps necessary in a trial including this last jump. To help participants quickly recognize the *dragging* option, and to match the visual character of the tapping side of the screen as much as possible, a continuous line (3 px in width) of the same dark red color connected the edge of the starting octagon and the target object.

The two paths leading to the objects were always separated by a right angle, with 45° angles opening from the axis through the starting octagon, perpendicular to the long side of the screen. Consequently, across trials, the two target objects were located in different positions along the same two lines originating from the middle of the starting octagon. Within one trial and within one participant, however, the two objects were located equidistantly from the starting location: The distances covered by either tapping or dragging were equal.

In addition to the objects described above, in one of the two individual conditions, the participants saw an instruction on the screen (Figure 4.1a). This text instructed the participants to collect the object either on the LEFT or RIGHT side of the screen. The text (48 pt, all capital letters) was centered along the x axis and located roughly in line with the target objects along the y axis (the bottom of the letters was at [$y_{ObjectCenter}$ - 15 px]). This was implemented to minimize the amount of visual scanning necessary to inspect the two targets in the layout and the instruction at the beginning of a trial.

4.2.4 Design

4.2.4.1 Action types and factors. We manipulated the distances to be covered on the screen with the collector octagon (Path factor), and the number of times that participants had to tap on the screen to move the octagon towards a target object (Step Number factor). The three levels of *Path* employed were 200 px, 400 px, and 600 px lengths; and the three levels of *Step Number* were 3, 5, and 7 steps to reach a target. All combinations of the three levels of each factor were used throughout the experiment. These were used in two individual conditions first, followed by a joint condition. In the latter condition, we used all possible combinations of the factor levels *within* and *between* the co-actors.

4.2.4.2 Conditions. The study had the following within-subjects conditions, in fixed temporal order:

1. Individual, No Choice condition. Participants were instructed to collect objects on a given side of the screen in each trial (left or right), reachable by tapping or dragging an equal amount of times. Movement times were collected in trials using every combination of the Path and Step Number factors. In total, 36 trials were presented to the participants in random order: 3 (Path) x 3 (Step Number) x 2 (mirroring the location of the tapping action – Left/Right side) x 2 repetitions of each trial. The shape of the black target object to be collected by tapping was randomized on every trial, as if by flipping a coin.

2. Individual, Choice condition. An individual 2AFC condition served to establish a baseline estimate of how costly each participant considered tapping compared to dragging. The 36 trials from the Individual No Choice condition were repeated in random order, without the instructions to pick a given side's object. We registered the type of action that participants chose in each trial.

3. Joint Choice condition. Following the two individual conditions, the participants completed a joint object matching task. As Figure 4.1b illustrates, both participants were presented with a tapping and dragging action option each. In every trial, Actor 1 decided that the dyad would collect either a black square pair or a black circle pair of objects. The identical shapes were positioned on the same side of the screen (left or right), and each pair could be collected by executing a combination of tapping and dragging distributed across co-actors. That is, if the participants wished to take into account the cost of the partner, they were required to compare tapping and dragging against each other both within their own task areas and between the two task areas of the two actors, since no object pairs could be collected by two tapping or two dragging actions.

In total, dyads completed 162 trials: 3 (Path for Actor 1) x 3 (Path for Actor 2) x 3 (Step Number for Actor 1) x 3 (Step Number for Actor 2) x 2 (the identity of Actor 1). Both actors completed the same 81 trials in randomized order, with the only constraint that they took turns in adopting the role of Actor 1. The shape of the object to be collected by tapping, and the side on which the tapping action appeared for Actor 1 were also randomized on a trial-by-trial basis.

4.2.4.3 Cost Disparities. We consider that action costs are monotonic functions of the path lengths that people had to cover on the screen by dragging their hands, or by tapping repeatedly on the screen to move the green octagon along a path to the target object. Since tapping costs were determined using Step Numbers, for comparability to the alternative dragging actions, we transformed step-based tapping costs into path lengths (i.e., functional distances) using individual PSE estimates collected in the Individual Choice condition. For example, let us say that executing 1 tap is equivalent in terms of costliness to dragging the octagon over a 120 px distance (i.e., PSE = 120 px/tap). In a trial such as the one in Figure 4.1a, where a participant had to decide between dragging the green octagon over a 600 px-long path on the right side or tapping 7 times over the same path length on the left side, using this equivalence, we would estimate that tapping would cost approximately 7 * 120 = 840 px. Such a transformation allowed us to calculate cost disparities on the pixel scale.

The cost of choosing the black square (A1 in Figure 4.1a) in a trial is the length of path to be covered to it either by dragging or tapping, depending on what kind of action had to be taken to reach A1 in the given trial. If the actor makes her decision between the two possible target objects based on her expected action cost, she should compare the cost (a1) of reaching A1 to the cost (b1) of reaching the object B1 (black circle). Self Disparity will therefore be determined as a1-b1, and similarly, Other Disparity is determined as a2-b2. Note that because of the complementarity of the actions to be executed by the co-actors, when Self Disparity is calculated as the difference between Actor 1's tapping and dragging costs (a1-b1 = $Tap_{Actor1} – Drag_{Actor1}$), then Other Disparity will necessarily be calculated as the difference between Actor 2's cost of dragging and tapping costs (a2-b2 = $Drag_{Actor2} – Tap_{Actor2}$), and vice versa. The joint cost of an action is again taken to be the summed costs of the co-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 is the sum of the two individual disparities (Self Disparity and Other Disparity).

In the example joint arrangement (Figure 4.1b), if we assume that both actors have the same PSE of 120 px/tap for transforming the Step Number into pixels, we can estimate the Self, Other, and Joint cost disparities on the same scale. For paths of length 400 px for Actor 1, Self Disparity will be (400 px – 3*120 px) = 40 px, and for paths of length 600 px for Actor 2, Other Disparity will be (7*120 px – 600 px) = 240 px. The positive signs of both disparities signal that the costs of collecting A1 and A2 (the two squares) are relatively costlier than collecting the circles, and therefore for both participants, it would be the individually efficient choice to collect the black circles. Further, the Joint Disparity in this example is 280 px, equal to (400 px + 7*120 px) – (3*120 px + 600 px). The positive Joint Disparity suggests that collecting the black circle pair (B1 + B2) is the co-efficient option in this trial, associated with overall lower action costs than the square object pair. This holds provided that the co-actors can be characterized as having the same points of subjective equality between the costs of tapping and dragging (i.e., the same PSE). However, we expected people to have a variety of "exchange rates" when they compare the subjective difficulties of tapping and dragging actions to each other. Therefore, we estimated participant-wise PSE values instead of a group-level PSE in the Individual Choice condition.

4.2.5 Procedure

Pairs of participants entered the lab at the same time and after giving their written informed consent, each participant first completed the individual conditions separately, sequentially. After the individual conditions, they completed the joint condition together. The procedures of the three conditions are described together below, as they were similar to one another; the differences are noted where necessary.

Before the individual object collection tasks, the participants read a list of step-by-step instructions on how to complete a trial. In the two individual conditions, they were instructed to collect a black object in each round of the game, with no requirement to complete the trials quickly.

They then completed a short practice session before the Individual No Choice condition to familiarize themselves with the task and the use of the screen and response keys. At the beginning of the joint condition, dyads also read step-by-step instructions before a short practice session. They were instructed to collect matching pairs of shapes by working together with their partner, without communicating with each other. The participants were not required to complete the task quickly and movement time was never mentioned in the instructions.

The steps of collecting objects were generally identical across all three conditions, as described in the following section. The difference between the two individual conditions was only that in the No Choice condition, the participants were required to follow the text instructions on the screen: They could only collect the instructed objects (located on the left or right side of the screen). In the Individual Choice condition, the participants were free to decide which object to collect. The joint condition differed from the individual conditions in the fact the participants performed a joint action sequence to collect object pairs, rather than individual objects: Actor 1 had a choice (like in the Individual Choice condition), but Actor 2 did not (like in the Individual No Choice condition).

4.2.5.1 Trial-by-trial procedure. In all three conditions, the participants were instructed to keep their dominant index fingers on the middle key of their response boxes to trigger the start of each trial. First, an orange-colored octagon appeared inside the starting location, which, in the joint condition, identified the participant who was required to start the trial (Actor 1). The participants in all three conditions were instructed to inspect the layout while the octagon was orange-colored, and in the two Choice conditions, to decide which target object they would pick up when prompted to move. In the Individual No Choice condition, the side instruction was displayed at this time already.

The octagon turned green after three seconds, to signal that the actor (in the joint condition: Actor 1) could start to move the octagon to collect one of the black objects. The first step of object collection was signaling the choice that the participant made. After releasing the middle key on the response box when the green octagon appeared, the participants were instructed to first press the key on the response box that was located closer to the side of the chosen object. That is, if participants chose the object on the left (right), they were required to press the left (right) key on the response box. This action turned the chosen black target object blue, to aid the participants by highlighting their current goal. In the joint condition, only Actor 1 had to perform this action, Actor 2's display showed the blue target highlighted as soon as it was chosen by Actor 1. Actor 2 could start moving his green octagon immediately after Actor 1 finished her part of the sequence.

If the participants chose the object reachable by dragging, they dragged the green octagon onto the (now blue-highlighted) target object with their index fingers. If they chose the object reachable by tapping, the participants were required to tap in the black-bordered octagonal starting area as many times as many dark red dots were displayed between the starting position and the target object, plus the final jump between the last dot and the target. An audio signal was played when the actors reached the object with the green octagon. They were not required to move the object back to the starting position. The trial was over in the individual conditions when the participant reached this stage. In the joint condition, once Actor 1 collected one of the objects, she pressed the middle key on her response box again to make the white octagon in front of Actor 2 turn green. The appearance of this second green octagon cued Actor 2 to start collecting the matching object on his side of the screen by performing the action complementary to Actor 1's action. The trial was over when Actor 2 collected the object with the shape corresponding to the one chosen by Actor 1. Non-matching objects could not be collected.

The participants completed the Individual No Choice condition in $M = 4.55$ minutes ($SD = 0.50$) and took on average $M = 4.48$ minutes ($SD = 0.60$) to complete the Individual Choice condition. The joint condition was completed on average in $M = 33.95$ minutes ($SD = 2.06$). Following the joint task, the participants responded to a short questionnaire on what they thought to be the purpose of the study and how much they liked their partner, using a 7-point Likert scale (1 – Not at all, 7 – Very much) before being debriefed about the experiment.

4.2.6 Data analysis

To test the hypothesis that the participants' object choices would be influenced by weighted additive Joint Disparities, we used hierarchical logistic regression models in a Bayesian parameter estimation framework. We fitted and compared three models in which the probability of choosing object A1 (the square) was predicted in turn by (1) *Self Cost Disparity*, (2) *Other Cost Disparity*, and (3) a weighted linear combination of the *Self and Other Cost Disparities*.

4.2.6.1 Cost transformation to the common scale. As mentioned in section 4.2.4 Design3 on the calculation of cost disparities across different types of action, we first needed to estimate individual points of subjective equality between tapping and dragging. The PSE was used as a measure of the "exchange rate" between the two actions. This was achieved by fitting psychometric curves to the decision data in the Individual Choice condition. We predicted the probability of tapping as a function of how much displacement (in pixels) one tap provided in a given trial (Tap Gain = Path / Step Number). To do this, we used generalized linear mixed models (GLMM) fitted with the glmer function from the *lme4* R package (Bates et al., 2015). Generalized linear mixed models are useful tools in estimating psychometric curves, especially if one is interested in individual-level estimates, beyond group performance. Following Knoblauch and Maloney's (2012) tutorial, we estimated thresholds for psychometric curves by fitting Gaussian cumulative distribution functions to the choice data (a probit model using the glmer function). The thresholds were calculated using the estimated beta coefficients and intercepts to yield the parameters of the quantile function[21]. The quantile function outputs the location on Tap Gain scale (the x axis of the psychometric curves) where the probability of tapping is 50%, i.e., the PSE (Knoblauch & Maloney, 2012).

The individual PSE values were used to transform the tapping costs, originally expressed in Step Numbers, into pixels for the Joint condition. These distances are equivalent to the functional

[21] The means and standard deviations of the Gaussian quantile function were calculated as: $\mu = -\beta_0/\beta_1$, and $\sigma = 1/\beta_1$, respectively (Knoblauch & Maloney, 2012, p.15.), where β_0 was each participant's intercept and β_1 their slopes estimated by the mixed-effects model with formula p(Tapping) ~ Tap Gain + (Tap Gain | participant).

distance defined by Rosenbaum and colleagues (2011). The transformation allowed us to follow the same Bayesian data analysis strategy as in Study 2 (Chapter 3). We fitted the logistic regression models mentioned above to the data of those participants that had PSE values within the range of the employed Tap Gain values in the experiment. That is, if a participant's PSE was outside of the range, for example, below 0, we did not observe a switch between behaviors along the Tap Gain dimension (see for an example participant 15's curve in Appendix A, Figure A.2). These participants' data were not analyzed, as their decision-making strategy could not be proven to be aligned with a cost minimization goal in the Individual conditions. To check if participants were trying to minimize movement durations in the Individual Choice condition, we first analyzed movement durations with a mixed model in the Individual No Choice condition to characterize the relationship between the time measure and Tap Gains (section 4.3.1 Individual No Choice condition2). We derived thresholds where switching between dragging and tapping would be beneficial for movement duration, as predictions for optimal choices. These thresholds were compared to the PSE in the Individual Choice condition (section 4.3.2 Individual Choice condition3).

Other Cost Disparity from the perspective of Actor 1 was determined in two ways: by transforming tapping costs into pixels using Actor 2's PSE value (which reflects the *true* costs and their disparities), and by using Actor 1's own PSE for the transformation (the *egocentric* Other Disparity). We decided to test the potential importance of egocentrically derived parameters in decision-making, since we cannot exclude the possibility that making inferences about a co-actor's cost function may be so cognitively demanding that actors would rather project their own cost functions onto their partner instead, to approximate their costs. Note that being able to calculate true Other Disparities relied on the availability of PSE values for both co-actors within a dyad. Therefore, in a subset of the sample where one of the two actors did not have a valid PSE, only the egocentric Other Disparity and Self and Other Disparity models could be fitted, not the true ones.

The dependent variable of all logistic regression models was the probability of Actor 1 choosing object A1 (the black square), and the cost disparity parameters served as predictors. The structure of the hierarchical model was the same as in Study 2, see Figure A.15. The value of the intercept in the logistic equation was fixed to 0 in all models, which is equivalent to assuming random decisions in the absence of any action cost disparities. As previously, the individual β coefficients were assumed to be normally distributed at the group level, corresponding to the assumption that participants' individual weighing strategies are noisy versions of a shared group-level weighing pattern. The weakly informed priors for this group-level distribution were set by hyperparameters $\mu \sim \mathcal{N}(0, 5)$, $\sigma \sim \mathcal{U}(0.0, 0.01)$, a wide distribution around a zero effect of cost disparity. The same hyperprior was used for all cost disparities. The individual and group-level posterior distributions of the beta coefficients were simultaneously estimated via Markov Chain Monte Carlo simulation in JAGS (Plummer, 2003), with the *runjags* package in R (Denwood, 2016) following Kruschke (2015).

To test the effects of further task parameters, the participants' proportions of object choices that were 1) collected by tapping, 2) co-efficient, 3) square-shaped, or 4) positioned on the screen's right side were tested against chance level using Wilcoxon signed-rank tests (two-tailed). As a measure of effect size, rank-biserial correlations are reported (Kerby, 2014). We report *V* statistics for the Wilcoxon tests, which correspond to the sum of ranks assigned to positive-signed differences between the tested paired samples and represents the value to be compared to those found in tables for Wilcoxon test.

4.3 Results

4.3.1 Individual No Choice condition

4.3.1.1 Movement durations as a function of Action Type, Path, and Step Number. The Individual No Choice condition served to determine how manipulations of the Path and Step Number factors and the type of action performed influenced movement duration. Movement

Duration was defined as the time elapsed between the moment the participant started to move the green octagon, and when the octagon reached a target object. In general, action type did not significantly influence the movement durations collapsed over different path lengths and step numbers (Figure 4.2, Taps: M = 0.887 s, SD = 0.26; Drags: M = 0.835 s, SD = 0.27; estimated coefficient of the linear mixed-effects model predicting Movement Duration by Action Type[22] only: β = 1.069, 95% Confidence Interval = [0.97, 1.18], p = .173).

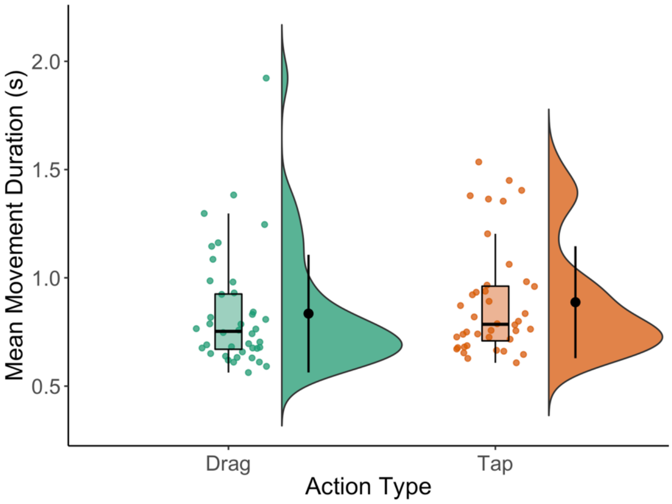

Figure 4.2: Mean movement durations (N = 40) of tapping and dragging in the Individual No Choice condition, collapsed across different values of Path and Step Number. In each box plot, the black horizontal line indicates the median, whiskers extend 1.5 times the interquartile range; the black dots signal the means, with their whiskers extending to one SD; and the density plots illustrate the distributions of mean movement durations. Each colored point represents a participant's mean according to the action type.

Next, we analyzed the Movement Duration values as a function of the *Path* and *Step Number* factors in a linear mixed-effects model[23] (details are reported in Table A.2). The main effects of Action Type and Path on durations were statistically significant (Action Type: β = 33.801, 95% CI = [1.71, 667.5], p = .021; log(Path): β = 2.417, 95% CI = [1.67, 3.51], p < .001). In addition, we found that the Action Type X log(Path) interaction effect was also statistically significant (β = 0.411, 95% CI = [0.25, 0.68], p < .001). These results suggest that for dragging actions, trials with longer paths lasted longer than trials with short paths – an effect which did not appear for tapping

[22] To examine movement durations, we ran three linear mixed-effects models with inverse Gaussian error distributions using a log link function, because the dependent variables were non-normally distributed. We report exponentiated β coefficients and their 95% confidence intervals here and below. Details of the models are available in Appendix A in Table A.1-A.4.
[23] The predictors were log-transformed to place them on similar scales and thus aid model convergence.

actions (**Figure 4.3**). These results indicate that the experimental manipulations worked: the action cost (estimated action duration) was the monotonic function of the relevant parameter (Path or Step) while it was not dependent on the irrelevant one (Path or Step).

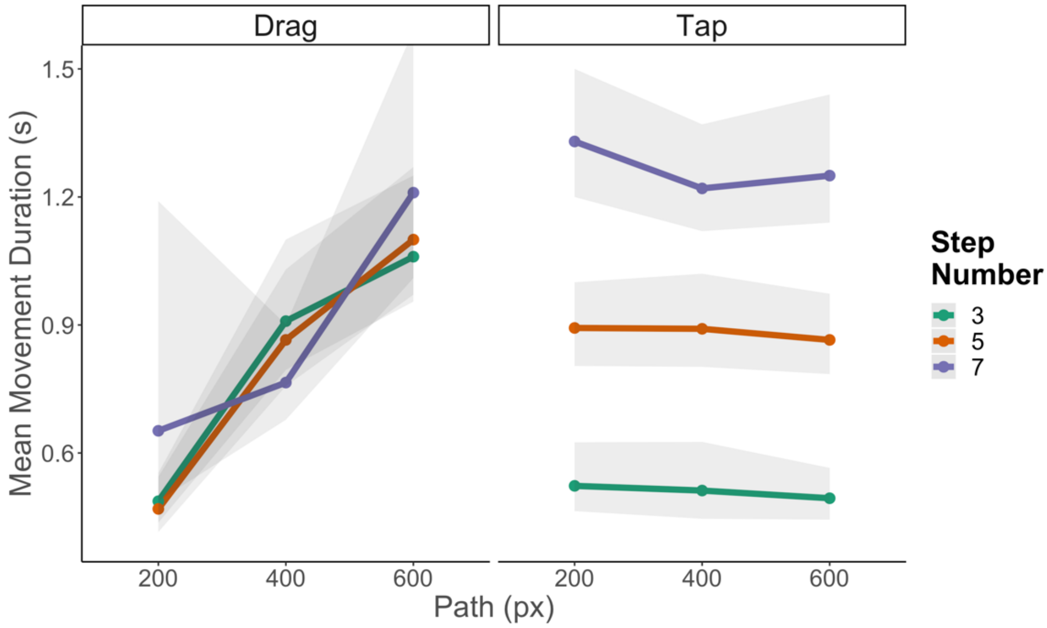

Figure 4.3: Mean movement durations in each combination of Path length (200, 400, 600 px) and Step Number levels (3, 5, 7) for the two action types in the Individual No Choice condition (N = 40). The grey shaded areas indicate bootstrapped 95% confidence intervals based on a 10000-step simulation.

4.3.1.2 Relative Movement Durations as a function of Tap Gain. To yield predictions for action optimization in the Individual Choice condition, we examined whether and how participants' relative movement duration times changed as a function of the experimental manipulations. Specifically, we calculated a Movement Duration Index as a measure of relative movement times when tapping and when dragging, by calculating the Movement Duration$_{Drag}$ / Movement Duration$_{Tap}$ ratio. When this index was > 1, tapping resulted in shorter movement times than dragging, and vice versa if the index's value was < 1.

We investigated if the Movement Duration Index was influenced by the experimental manipulations of the Path and Step Number factors by fitting a linear mixed model with the predictor Tap Gain (Table A.4). As described in section 4.2.6 Data analysis, the Tap Gain measure expresses the magnitude of octagon displacement as a result of one tap, which places all combinations of the two experimental factors on the same scale. We found a statistically significant positive relationship between the Tap Gain of trials and the Movement Duration Indices (β = 1.013, 95% CI = [1.011, 1.014], $p < .001$). The closer one tap moved the octagon towards a target object, the more beneficial tapping was to movement time minimization over dragging (Figure 4.4). A one px/tap increase in Tap Gain was expected to result in a 1.3% (95% CI = [1.1, 1.4]) increase in the relative movement duration of tapping over dragging, which for a 1 cm on-screen increase in Tap Gain translates into a 27% (95% CI = [23.9, 30.3]) increase in the Movement Duration Index.

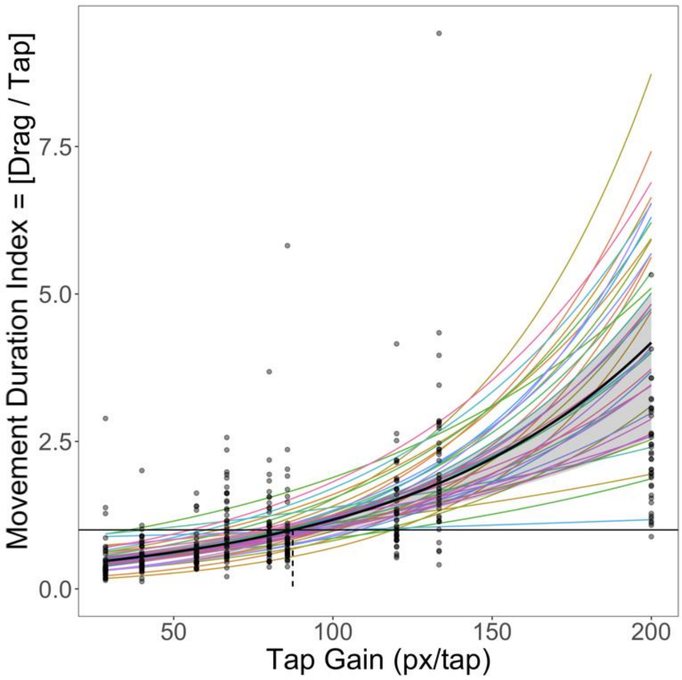

Figure 4.4: Relationship between Tap Gain and Movement Duration Indices in the Individual No Choice condition. Fitted linear regression curves are shown for each participant in color, the black curve shows the estimated group-level fixed effect of Tap Gain. The dark grey shaded area indicates a bootstrapped 95% confidence interval based on a 200-step simulation. Black points show individual data (N = 40): mean Movement Duration indices for each of the 9 levels of Tap Gain. The vertical dashed line indicates the estimated group threshold at 87.38 (calculated using the group-level fixed-effect estimate for Tap Gain), the level of Tap Gain where dragging and tapping took equal amounts of time.

On the individual level, almost all participants' relative movement durations reflected this general pattern: Above a given threshold value of Tap Gain (M = 86.23, SD = 18.4), participants tended to complete trials faster by tapping than dragging, (Figure A.1). This provides an objective benchmark for action optimization under the current task constraints, with the assumption that people aimed to optimize their actions for time. We investigated if participants chose their actions according to such an optimization principle in the Individual Choice condition.

4.3.2 Individual Choice condition

4.3.2.1 Descriptive statistics. On average, participants chose objects reachable by tapping actions in 54% (*SD* = 0.19) of the trials in the Individual Choice condition. This proportion was not statistically significantly different from chance (Wilcoxon signed-rank test[24] of the proportion data to chance level at .5: $V = 427$, $p = .067$, 95% CI for proportion .54 = [.49, .62], $r = -.04$). Similarly, participants chose the black square object (A1) on average in 53% of the trials (*SD* = 0.13; comparison to chance level at .5: $V = 375$, $p = .511$, 95% CI for proportion .51 = [.47, .56], $r = -.09$). Objects on the right screen side were chosen in 55% of the trials (*SD* = 0.14), which was a proportion different from chance (comparison to chance level at .5: $V = 418.5$, $p = .013$, 95% CI for proportion .54 = [.50, .60], $r = .02$). Overall, participants were not biased to choose either action type or target object over the other, and they exhibited a very small bias to choose objects on the right side of the screen.

4.3.2.2 Psychometric functions fitted to tapping decisions. To test whether the participants optimized their actions for the time spent moving the octagon, and to estimate their PSE to transform the costs of tapping into pixels, we analyzed choices in the Individual Choice condition as a function of Tap Gain. The probability of participants choosing objects reachable by tapping was in a statistically significant, positive relationship with the gain of a tap (Figure 4.5), based on a mixed-effects probit model ($\beta = 1.007$, 95% CI = [1.004, 1.010], $p < .001$, see Table A.5). That is, the farther the participant could move the octagon by one tap, the likelier it was that she chose the tapping action. A 1 px increase in the gain of a tap resulted in an expected 0.7% (95% CI = [0.04, 0.10]) increase in the z-score of the probability of choosing a tapping action, that is, a 1 cm increase in Tap Gain was associated with a 15% (95% CI = [8.5, 21.8]) increase in the probability of choosing to tap instead of drag.

[24] To correct for multiple comparisons, the α-levels for the Wilcoxon tests reported here and in sections 4.3.3 Joint Choice Condition1 and 4.3.6 Excluded Participants were divided by 3 ($\alpha = .017$).

4.3.2.3 Individual points of subjective equality (PSE) estimates. The PSE, or threshold Tap Gain value where dragging was equally likely to occur as tapping, was calculated for each participant using their estimated random intercepts and slopes (Figure A.2; Knoblauch & Maloney, 2012). It was not possible to calculate the PSE for those participants whose fitted curves' slopes were negative, i.e., a cumulative normal distribution function could not be fitted to their decisions (n = 5). The decision thresholds (n = 35, M = 96.11, SD = 243.26) were compared to the relative movement duration thresholds estimated in the Individual No Choice condition (n = 40) with a Bayesian paired-samples t-test[25]. A difference between the two thresholds would suggest that participants did not optimize their actions for the time spent on moving the octagon in a trial. We found that the data were likelier under the null hypothesis of the two threshold distributions being similar, than the alternative hypothesis of a difference (n = 35, BF_{01} = 5.37, 95% credible interval for the difference = [-0.353, 0.279], median effect size: -0.036, Figure A.3). This suggests that overall, the participants to whose data psychometric curves were successfully fitted (n = 35) made decisions in the Individual Choice condition that likely optimized their movement duration time. A visual inspection of the individual curves in Figure A.2 suggests that this finding was driven by at most 26 out of 35 participants (see also section 4.3.4 Parameter Estimations1).

[25] The Bayesian t-test was run in JASP (JASP Team, 2020), using the default Cauchy prior with width 0.707.

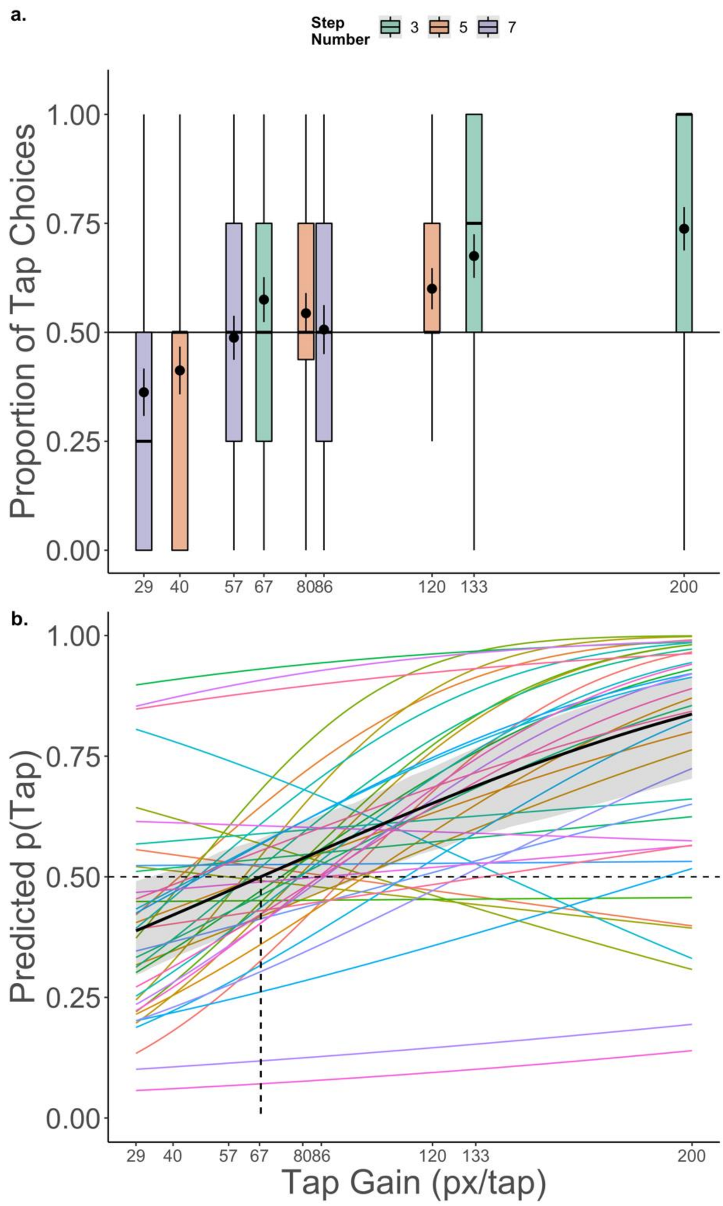

Figure 4.5: (a) Observed mean tap choices as a function of Tap Gain in the Individual Choice condition (N = 40). In each box plot, the black horizontal lines indicate medians, whiskers extend 1.5 times the interquartile range, and the black dots signal means, with whiskers extending to one SD. Step Numbers are given to indicate which combinations of Path X Step Number elicited given ratios of tap choices. (b) Predicted probabilities of tapping as a function of Tap Gain in the Individual Choice condition. Fitted probit regression curves are shown for each participant (N = 40) in color, the black line shows the estimated group-level fixed effect of Tap Gain. The dark grey shaded area indicates a bootstrapped 95% confidence interval based on a 200-step simulation. The vertical dashed line indicates the estimated group threshold at 67.05, the level of Tap Gain where dragging and tapping were equally likely to be chosen. In both figures, the Tap Gain values are rounded to the next integer.

The estimated PSE values approximate the "exchange rates" at which each individual chose the two types of actions with equal probability. We used these values to estimate how costly people perceived tapping to be to move the octagon over a distance, compared to dragging the octagon over that same distance. Based on the relationship Path/Step Number = PSE, we estimated the cost of a tapping action in a given trial in pixels as $Cost_{Tap}$ = PSE * Step Number. These were used along with the expected costs of dragging to calculate the Self and Other cost disparities, as described in section 4.2.6 Data analysis1, which in turn were used to predict the probabilities of A1 object choices in the Joint Condition.

4.3.3 Joint Choice Condition

4.3.3.1 Descriptive statistics. To check for any biases, we calculated the overall proportions of square object choices, tapping, and right-side choices. On average, participants chose object A1 in half of the trials (M = 0.50, SD = 0.06), tapping actions in 48.9% of the trials (SD = 0.11), and objects on the right side of the screen in 52.5% of the trials (SD = 0.15). Neither of these proportions were statistically significantly different from chance (Wilcoxon signed-rank tests[26] to compare against chance level at 0.5; square choice: V = 424.5, p = .851, 95% CI for proportion .51 = [.48, .52], r = .04; tap choice: V = 333, p = .303, 95% CI for proportion .49 = [.45, .53], r = -.19; right side choice: V = 468.5, p = .436, 95% CI for proportion .52 = [.46, .57], r = .14).

4.3.4 Parameter Estimations

4.3.4.1 Analyses of sub-samples. As previously mentioned, the PSE of the fitted normal cdf in the Individual Choice condition could not be calculated for those participants for whom the relationship between the probability of tapping choices and Tap Gain was negative (n = 5). Further, we found that 9 participants' estimated PSE values were out of the range of Tap Gains used in the experiment (28.57-200 px/tap). We excluded these 14 participants from further analyses.

A. Complete dyad data. In 8 dyads, both co-actors' PSE values were in the valid range, therefore, their action costs were transformed based on the two participants' respective cost exchange rates (both "true" and "egocentric" Other Disparity could be calculated, see section 4.2.6 Data analysis1). To the data of this subsample (n = 16), the following models could be fitted: (1) *Self Cost Disparity*, (2) true *Other Cost Disparity*, (3) egocentric *Other Cost Disparity*, (4) a weighted linear

[26] The Bonferroni-corrected α-level for the multiple comparisons was .05/3 = .017.

combination of the *Self and true Other Cost Disparities*, and (5) of the *Self and egocentric Other Cost Disparity*.

B. Partial dyad data. For an additional 10 participants, whose PSEs were calculated, their partners' data were excluded. The transformation of these participants' partners' tapping costs into pixels was based on the 10 participants' own PSE values (i.e., only "egocentric" Other Disparities were calculated). We fitted three Bayesian models to their data: (1) *Self Cost Disparity*, (2) egocentric *Other Cost Disparity*, and (3) a weighted linear combination of the *Self and egocentric Other Cost Disparity*.

The results of the parameter estimations for these two subsamples were highly consistent. In both cases, the models combining Self and egocentric Other Disparity fit the data best. Due to the similarity of the results, we pooled the data of the two subsamples (n = 26) and report results of the Bayesian parameter estimation for the (1) *Self Cost Disparity*, (2) egocentric *Other Cost Disparity*, and (3) *Self and egocentric Other Cost Disparity* models on the pooled data. The joint distributions of these disparity parameters in the pooled subsample are shown in Figure 4.6. Details of the other two subsamples' results are available in the Appendix (Table A.7 -8, Figure A.6-14).

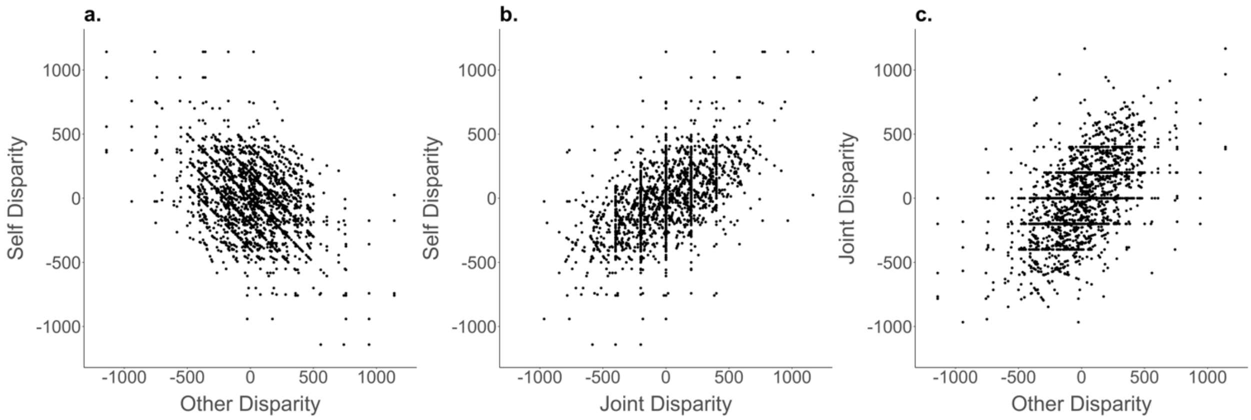

Figure 4.6: Scatterplots of the joint distributions of cost disparities, collapsed across all trials of the pooled subsample's dyads (n = 26, 2106 trials). (a) Self and Other Disparities were negatively correlated, r = -.366. (b) Self and Joint Disparities, and (c) Joint and Other Disparities were positively correlated with each other (both r = .563).

4.3.4.2 Co-efficient object choices in the pooled subsample. Based on the egocentric Joint Disparity values calculated in the pooled subsample (n = 26), we first examined object choices coded as co-efficient. We found that participants chose the co-efficient object on average in 69.4% (SD = 0.10) of the trials in the Joint condition. To test whether this proportion was higher than the proportion we would expect if participants made self-cost minimizing decisions, we compared the 69.4% to chance level at the proportion of trials in the subsample in which co-efficiency and self-cost minimization predicted the same object choice (i.e., congruent trials, see Chapter 2). The proportion of co-efficient choices was only marginally higher than chance[27] so defined, which was a moderate effect (Wilcoxon signed-rank test to .65: $V = 252$, $p = .053$, 95% CI for proportion .70 = [.64, .75], $r = .44$).

The prior predictions of strategies aiming to minimize the Self-, Other- and the weighted combination of Self and Other action costs for this pooled subsample are shown in Figure 4.7d-f, with the observed A1 choices in Figure 4.7a. In the text and figures of parameter estimates we report rescaled coefficients that express the effect of the cost disparities in units of one on-screen cm instead of one pixel. Pixel-based estimates are reported in the Appendix (Table A.6).

[27] The proportion of co-efficient choices was not related to how much participants rated liking their co-actor in the post-task questionnaire (*Mdn* = 6, interquartile range, *IQR* = 1.75; Spearman's $\rho = .140$, $p = .494$).

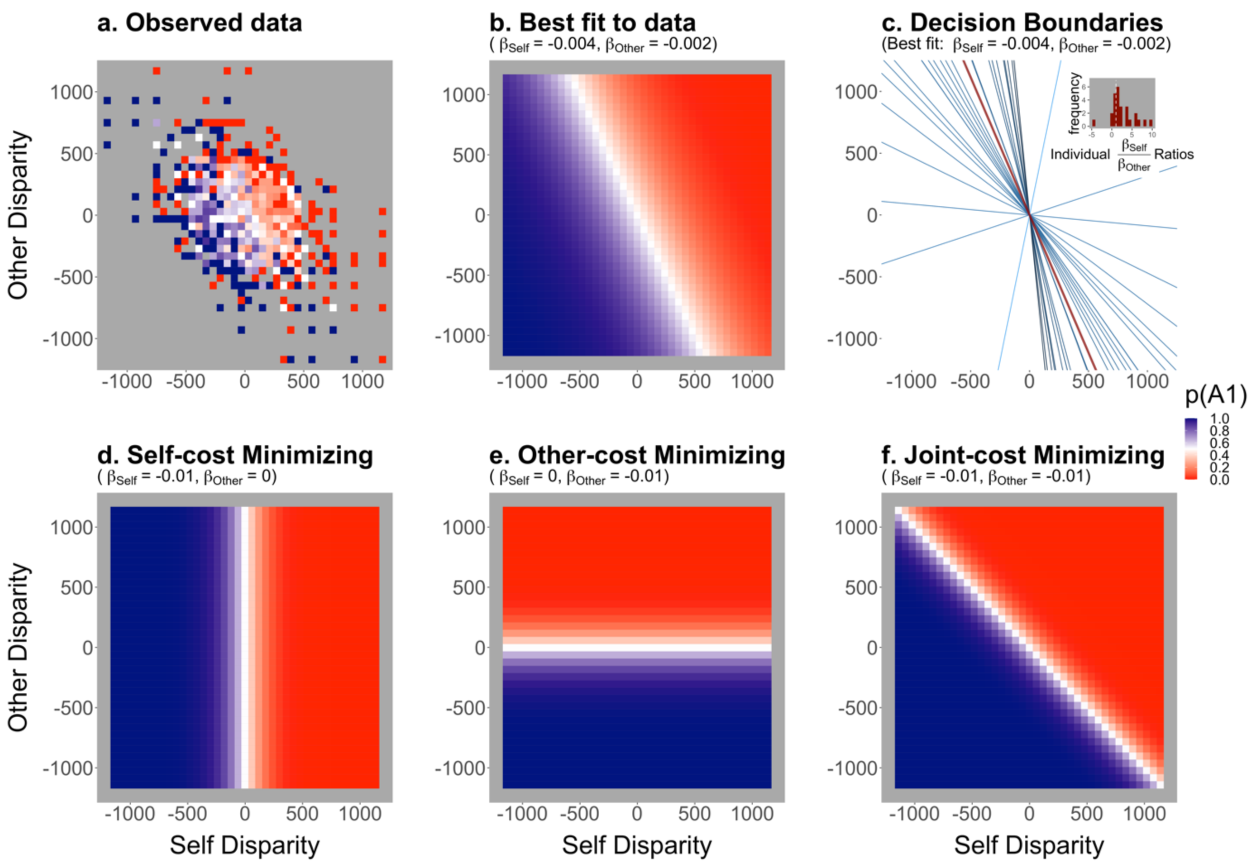

Figure 4.7: (a) Observed object A1 choices (n = 26, bin width = 60), and (b) the posterior predictions of the best-fitting model using the linear combination of Self and egocentric Other Disparities. (c) Individual decision boundaries according to the best-fitting model; inset: frequency distribution of the $\mu_{\beta Self}/\mu_{\beta Other}$ ratios (see also Figure A.4). The vertical white dashed line denotes 1, equal weights on Self and Other Disparity. (d-f) Predictions for optimal responses according to Self, Other, and Joint (i.e., Self + egocentric Other) cost-minimizing strategies, respectively. The lower the disparity to be minimized according to a model, the higher the probability of picking object A1 (blue). Predictions were calculated assuming that one pixel increase in a given parameter would result in 1% decrease in the odds of choosing object A1 over B1. All plots feature disparities in pixels.

Model 1: Self Disparity. The results of the parameter estimation suggest that by itself, Actor 1's own cost disparities significantly influenced the probability of picking object A1 in the predicted direction. The mode of the posterior for $\mu_{\beta Self}$, the parameter denoting the group-level weight for the cost disparity, was -0.071, and the distribution's 95% Highest Density Interval was entirely below zero (Figure 4.8a, 95% HDI = [-0.102, -0.041]). Based on this model, a 1 cm increase in Self Disparity was expected to result in a 6.8% decrease in the odds of picking object A1 over B1.

Model 2: Other Disparity. Actor 2's cost disparity, estimated based on Actor 1's own PSE, had a similar effect on the probability of square object choices in the predicted direction. The mode of the posterior distribution for $\mu_{\beta Other}$ was -0.016 (Figure 4.8b, 95% HDI = [-0.029, -0.005]), and the majority of the mass of the distribution was below zero. An increase in Other Disparity by 1 cm is expected to lead to a 1.6% decrease in the odds of an A1 choice over a B1 choice.

Model 3: Self and Other Disparity. The group-level means ($\mu_{\beta Self}$ and $\mu_{\beta Other}$) of the β_{Self} and β_{Other} coefficients for both disparities in Model 3 were distributed below zero (Figure 4.8c, Self: 95% HDI for $\mu_{\beta Self}$: [-0.118, -0.055], Mode $\mu_{\beta Self}$ = -0.084; Other: 95% HDI for $\mu_{\beta Other}$: [-0.054, -0.022], Mode $\mu_{\beta Other}$ = -0.038). Increasing Self and Other disparities each by a cm whilst holding the other disparity constant is expected to lead to an 8.1% and a 3.7% decrease in the odds of picking object A1 over B1. The relative average weights on Self and Other Disparity according to this combination model were .69 (95% HDI: [.45, .97]) and .31 (95% HDI: [.18, .44]), respectively.

4.3.5 Model Comparison

The three fitted models were compared based on Leave-one-out cross-validation (Gelman et al., 2014; Vehtari et al., 2017). The lower the LOO-CV measures for a given model, the better its expected accuracy in predicting future data. Here we only report LOO-CV, but Table A.6-8 also present WAIC values, as in Chapter 3 (Watanabe-Akaike Information Criterion, Gelman et al., 2014; Watanabe & Opper, 2010). We also fitted ROC curves (Fawcett, 2006) to each model's posterior predictions and the observed object choices and compared Area Under the Curve measures (Figure A.5). The best model according to all metrics was the model using the weighted combination of Self and Other Disparity (LOO-CV = 2400.5 [SE = 39.7], AUC = 70.9%). Figure 4.7b shows the posterior predicted probabilities of A1 choices based on this model, with individual decision boundaries illustrated in Figure 4.7c. The second-best model was the one including Self Disparity by itself (LOO-CV = 2528.4 [SE = 34.3], AUC = 67.3%), with the minimization of Other costs being the least likely explanation of decision data (Other Disparity model: LOO-CV = 2871.3 [SE = 14.1], AUC = 55.7%).

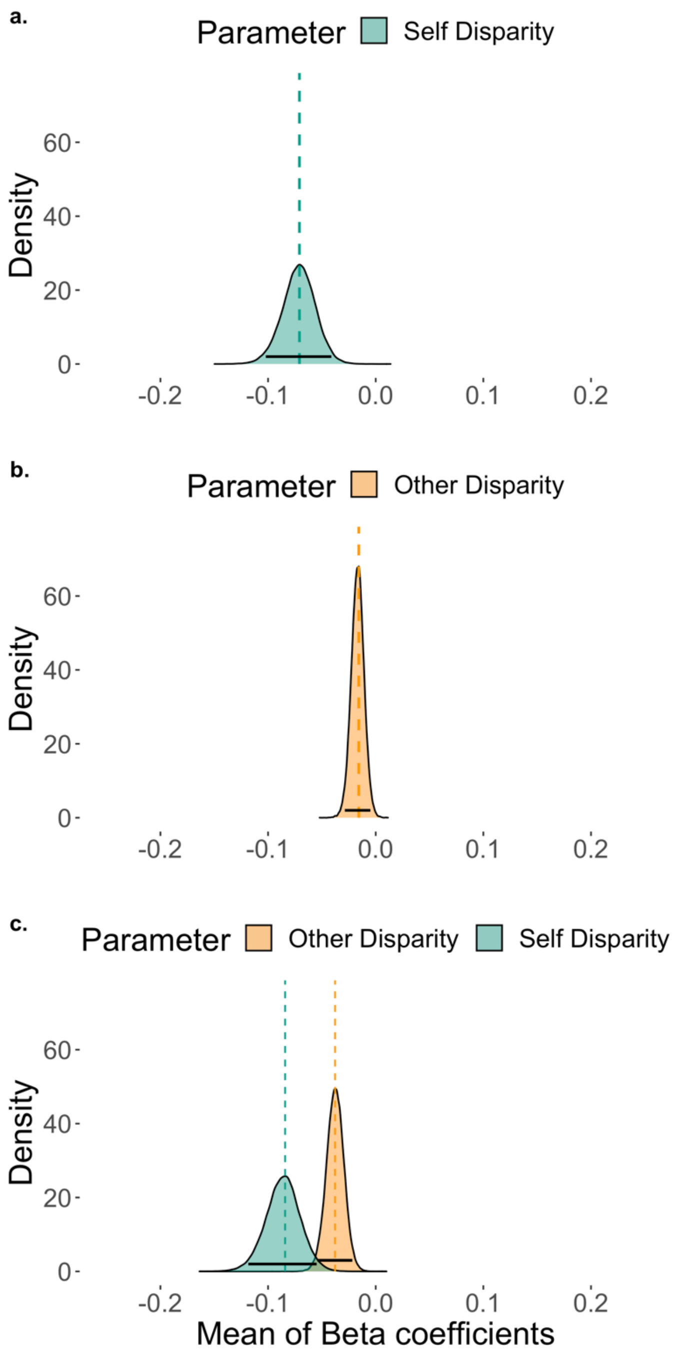

Figure 4.8: Posterior probability distributions of the rescaled μ_β parameters for Self and Other Disparities in the (a-b) single predictor models (Models 1 & 2), and in the (c) combination model (Model 3). The beta coefficients are scaled to express the effect of cost disparities on object A1 choices in cm units. The dashed vertical lines indicate the Mode μ_β, the black horizontal lines represent the 95% highest density intervals of each distribution.

4.3.6 Excluded Participants

To characterize the behavior of the participants excluded from the parameter estimation in the Joint condition (n = 14), we briefly summarize their decisions in both the Individual and the Joint Choice conditions. They were excluded because they did not seem to have made decisions to minimize movement duration or some related variable. We concluded this either because their tapping action choices in the Individual Choice condition were negatively related to Tap Gains (n = 5), or because their estimated points of subjective equality were outside the range used in the experiment (lower than the minimum: n = 6, higher than the maximum: n = 3).

In the Individual Choice condition, these participants seemed to make decisions not different from chance: they chose the square on average in 52% of the trials (SD = 0.11; Wilcoxon signed-rank test[28] of the proportions to chance level at .5: V = 58.5, p = .728, 95% CI for proportion .51 = [.44, .61], r = .11), and chose the object on the right hand side of the screen on average in 56% of the trials (SD = 0.16; Wilcoxon signed-rank test of proportions to chance level at .5: V = 36, p = .412, 95% CI for proportion .53 = [.46, .72], r = -.31). On average, they chose tapping actions in 57% of the trials (SD = 0.29; comparison to .5: V = 64, p = .207, 95% CI for proportion .57 = [.31, .81], r = .22).

In the Joint condition, we similarly found that on average, these participants were not biased to choose one target object over the other (square choice proportions: M = .53, SD = 0.07; comparison to chance level at 0.5: V = 70, p = .285, 95% CI for proportion .52 = [.48, .58], r = .33), to choose tapping over dragging (tapping proportions: M = .47, SD = 0.13; V = 38.5, p = .396, 95% CI for proportion .48 = [.36, .58], r = -.27), or to choose the object on the right hand side (M = .52, SD = 0.11; V = 66, p = .414, 95% CI for proportion .52 = [.44, .60], r = .26).

4.4 Discussion

In the present study, we examined the question how people plan joint action sequences, when planning co-efficient actions would require the integration of costs of different action types. We hypothesized that to the extent that in individual settings actors are able to select actions that make them relatively efficient, in joint actions they would integrate their partner's estimated relative costs into the planning process to plan co-efficient action sequences.

[28] The Bonferroni-corrected α-level for the multiple comparisons was .05/3 = .017.

Participants collected objects on their own in an individual 2AFC task and collected object pairs together with a partner in a joint task. In both settings, they could choose between two means of object collection: tapping repeatedly on a touchscreen to move an octagon towards a target or dragging the octagon with one continuous movement onto the target. To act in an efficient manner, they needed to compare the relative costs of the available action options on each trial. The costs were operationalized as Path (distance to target) and Step Number (amount of taps required). In the joint condition, the decision-maker was faced not only with two potential actions for herself, but with two potential action sequences for the dyad, composed of different combinations of tapping and dragging. Therefore, in the joint task, cost comparison and minimization could happen not only with respect to the decision-maker (*Self-cost minimizing model*), but also to her partner (*Other-cost minimizing model*), and to the dyad (*Joint-cost minimizing model*), focusing on the relative joint costs of the combined tap-and-drag sequences.

First, our hypothesis regarding action choices in individual settings was confirmed. We found that 65% of the participants (26 out of 40) chose actions in the individual task in a consistent way that reflected that they tried to minimize a common currency behind actions, likely their movement duration times (i.e., the time spent on performing all required taps or one drag in a trial). For the majority of participants, there was a strong positive relationship between the gain of one tap and the relative time-benefit of tapping, as suggested by the Individual No Choice condition's results. That is, the larger the distance over which one tap moved the octagon on the screen, the faster a tapping trial was performed, relative to a dragging trial. This relationship implicitly set the rule for time-based action optimization, which we can roughly summarize as "*If one tap takes your octagon closer to the target at least by 87 pixels (~4.6 cm) on the screen, tap; otherwise drag*".

This simple rule seems to have been recognized and followed by 26 participants who tended to choose tapping over dragging consistently – more often as the gain of a tap increased – as suggested by their decisions in the Individual Choice condition. We found that the group's estimated "exchange rate" (equivalent to the weighing constant in Rosenbaum et al., 2011) between tapping and dragging was around 67, meaning that on average, tapping once was as costly as dragging over 67 pixels onscreen (approx. 3.56 cm).

Fourteen participants did not seem to follow a time-based optimization rule in the present task: they were either biased towards one action over the other, or tended to choose at chance, based on the analysis regarding choices and tap gains (see the individual decision plots in Figure A.2). We can speculate about various potential explanations for these participants' behavior. On the one hand, it is possible that some of them found the cost comparison too challenging and therefore resorted to some other decision rule, like always choosing the same action. On the other hand, it is possible that they could compare the relative costs of actions and preferred to keep tapping in the same area because they found the repeated extension movement of the arm in dragging actions throughout the task too effortful, and therefore costly (the maximum distance used was 600 pixels, ~31.8 cm). This would mean that they used a common currency other than time or functional distance. Lastly, perhaps they did not care about time, distance, or movement effort, and simply found one movement type less monotonous and boring than the other; or made random choices. Regardless of their exact strategies, we conclude that the majority of the participants made consistent decisions in the individual setting, and more than half of the sample (65%) acted to minimize some measure of cost, a hidden common currency, which may or may not have been time or time-related functional distance[29]. This finding aligns with the literature on judging the relative costs of different sensorimotor actions like reaching and walking, which has proposed functional distance as a reasonable proxy for the common currency (Rosenbaum, 2008; 2012; Rosenbaum & Gaydos, 2008; Rosenbaum et al., 2011).

Secondly, our main hypothesis concerning joint action planning was that to the extent that people make cost-minimizing decisions between actions of different types in an individual setting, they will also incorporate their partner's action costs into their joint action planning and choose

[29] N.B. The hidden common currency was at the minimum related to time so we could manipulate it in our experiment, suggested by the relationship we found between movement durations and the gain of one tap.

co-efficient sequences. We found support for this hypothesis based on the results of the logistic regression modelling that we conducted following the methodology of our previous study (Chapter 3), where we analyzed object choices in the joint task predicted by cost disparities. As described above, we found that 65% of the participants acted efficiently in the individual condition, so we focused on these participants' behavior in the Joint Choice condition. Based on each individual's estimated relative costs between tapping and dragging, we calculated first the functional distance of tapping, and then the cost disparities for Self and Other between tapping and dragging in each trial.

The results of the parameter estimations and model comparisons are in line with previous results (Chapter 3) and support our joint action planning hypothesis. On average, people chose the co-efficient action option in 69.4% of the trials, which was a proportion numerically larger than chance (.65), although it did not reach statistical significance. The decisions in the joint object collection task were best explained by a model that incorporated both the decision-maker's and the co-actor's relative action costs (*Self and Other Disparity model*). That is, the higher the combined costs of a dyad's tap-and-drag sequence, the less likely that it was the chosen course of action in a trial, which reflects a co-efficiency maximizing strategy at play. The relative average weights on Self and Other Disparity in the decision-making were .70 and .30, respectively. This reflects a tendency of people to assign more importance to their own individual efficiency, but overall, throughout repeated interactions, they took into account both co-actors' action costs. This pattern of findings is in general accordance with the co-representation and facilitation literature, especially with Ray et al. (2017) who showed that people are sensitive to the potential future efforts of a co-actor.

It is worth noting that this relative Self-oriented bias was not only reflected in the weights placed on the disparities, but in the fact that decisions were better explained by a model which used the egocentrically estimated cost disparities of a partner (using the decision-maker's own cost exchange rate), rather than those estimated based on the relative costs of the partner (see section 4.3.4 Parameter Estimations1). This finding suggests that even if people act efficiently on their own when comparing two different actions, and they take into account a co-actor's costs in a joint action, they may do so based on their *own* judged relative costs. It is possible that in cases of composite joint action sequences where there is uncertainty about the interaction partner's subjective costs, simulating the potential actions of the co-actor based on their perceptible task constraints (i.e., path lengths, step numbers on the screen) is the most computationally efficient solution for joint action planning. Ray et al. (2017) also found that after the initiator performed the task of her co-actor, on a subsequent joint task, she facilitated her co-actor's potential actions to a larger extent than before having experience with the other's task. The authors interpreted this as indication for the role of action simulation in representing a partner's future task difficulty and effort. Our findings regarding the better explanatory power of egocentrically estimated Other disparities are consistent with this view.

The experiment raises various questions and suggests potential avenues for future research. First of all, as we have emphasized throughout, we cannot conclude with certainty that the common currency behind judging the relative costs of tapping and dragging in the task was the duration of these actions. Functional distances are assumed to be related to these times, since the longer a physical distance, the longer it takes to cover it. Future work could explore in more detail what the common currency might be in a task like ours, for example, by a priori fixing the duration of a trial.

Secondly, the role of simulation in co-efficient planning would also be worth investigating in future studies. This could be addressed with a design similar to Ray et al. (2017), by introducing a new kind of action for Actor 2 in an initial joint choice task (an action which Actor 1 did not perform previously in the individual baseline conditions), then providing an opportunity for Actor 1 to try this new action in an intermediary individual condition, and re-testing with a second joint choice task. An increase in the weight applied to the second actor's individual action costs would inform us of the importance of simulation in the joint-cost integration process.

On the other hand, it would also be interesting to explore under which conditions people stop using simulation when they compare different action options in joint contexts. It might be possible that simulation is not used to compare potential joint action plans, as it might take too much time to make a decision this way. Based on reaction time data, Rosenbaum (2012) concluded that in his task, participants compared two action sequences by parallel search of their features rather than by sequentially simulating the entire walking-reaching sequences. One way to do this is that in a two-stage process, first a very low-cost option was accepted (or a very high-cost option rejected), and if such an option was not available, in a second stage, more detailed comparisons were made. Another possibility is that if one option was sufficiently less costly than the other, then it was chosen, and more time was spent on making comparisons the smaller the cost difference between the options (Rosenbaum, 2012). An investigation of the decision times in planning co-efficient joint actions could provide insight into what decision processes underlie the integration of joint costs – whether people simulate the full joint plan (i.e., every potential movement to the self and the partner, and their combinations) or run a parallel search to compare cues for relative costliness, without simulating the entire joint plan. For instance, we could test how the decision time patterns differ between tasks that employ only one type of action (e.g., the task in Chapter 3), and two types of action as in the present study. It is possible that a parallel comparison model would better account for decision-making in situations where the costs are readily comparable on the same scale of say, distance; whereas comparing composite joint action plans of two different actions would be likelier to engage serial simulation processes.

Last but not least, the investigation of joint action planning would take one further step towards ecological validity by addressing situations where a shared goal could be achieved by performing a combination of sensorimotor and cognitive tasks distributed across actors. For example, a group of friends having to assemble a roomful of IKEA furniture on a very tight schedule poses a challenge that involves multiple kinds of sub-tasks: reading instructions, organizing furniture pieces, clearing an area for work, distributing tasks among individuals in the group, assembling furniture. Real-life problems like this could be modelled on smaller scales using tasks similar to our present experiment, when extended to comparisons of mental and physical effort.

In summary, the present study aimed to extend the investigation of co-efficiency to more realistic coordination problems by modelling situations where different actions are weighed up against and combined with each other to reach a common goal. In such settings, planning a co-efficient joint action would require the integration of two different kinds of action costs, distributed across the co-actors. We found support for the hypothesis that conditional on individually efficient, or at least consistent[30], decision-making, people are able to integrate the relative costs of actions available to themselves and to their partner, and they make decisions that relatively minimize these joint costs. It is worth reiterating that we cannot draw strong conclusions about which common currency people used for cost-minimization, only that participants seemed to minimize a currency for which time seemed a suitable proxy. Furthermore, we could not directly compare trial durations and expended energetic costs between co-efficient and sub-efficient actions in the joint choice condition to test if the participants acted efficiently in the sense of energy-, biomechanical effort- or time-minimization. Nevertheless, our study provides evidence for the idea that, based on the estimated composite joint costs of a dyad in repeated interactions, people are able to make decisions that will minimize action costs relative to the alternatives, that is: they make collectively rational decisions.

[30] Consistent in the sense that action decisions monotonically changed with a variable, here, along the scale of Tap Gain.

Chapter 5. Co-efficiency as a potential focal point in coordination problems

5.1 Introduction

To reach shared outcomes in joint actions in a coordinated manner, interacting agents rely heavily on mechanisms such as action monitoring, prediction (Vesper et al., 2010), communication (verbal or sensorimotor, Pezzulo et al., 2019), and appealing to shared task representations (Vesper et al., 2017). Sometimes, however, the agents might have limited information about their partner's behavior, and communication might not be possible to reduce uncertainty. Such coordination problems may not be solved by most of the aforementioned mechanisms – yet anecdotal and experimental evidence suggest that people are fairly successful at coordinating decisions with remote interaction partners.

In *The Strategy of Conflict*, Schelling (1960) described his informal experiments where participants were told to coordinate with others by both naming "heads" or "tails" for a prize, or to imagine that they had to meet a friend somewhere in New York City but they could not negotiate the time and place of the rendezvous. Where and at what time would they meet? He found that in the first survey, the majority of respondents chose "heads" (36 out of 42) and in the second, most opted to meet their friend at Grand Central Station at noon. Schelling (1960) called the outcomes that people coordinated on *focal points,* the roles of which in pure coordination games were later confirmed in the lab (Mehta et al., 1994a, 1994b). Seemingly, limited information and a lack of communication does not preclude successful coordination as long as people can identify the same focal point to coordinate on. How do people succeed at this in joint actions? The present chapter explores the potential role of action co-efficiency as a focal point in coordination problems similar to the NYC meeting scenario: where reaching a coordinated state requires movements.

The New York City meeting problem has primarily been discussed as a pure coordination problem without conflicts of interest between interaction partners (Bacharach & Bernasconi, 1997;

Colman, 1997; Sugden, 1995), but if we consider it a joint action scenario and take into account the costs related to travelling to the agreed upon location, we will likely introduce payoff conflicts between the friends (cf. Lewis, 1969/2008)[31]. In order to coordinate with a co-actor in a joint action to reach a shared goal, individuals might need to take on tasks that are more or less effortful than their partner's. Such a situation might be better conceptualized as a mixed-motive or bargaining game (Schelling, 1960). In the present chapter, we address the questions whether joint-payoff related focal points help coordination, in spite of asymmetries in individual payoffs (here: action costs) between the co-actors.

Most of previous research addressed focal point-reasoning in a behavioral economics framework[32]. These experiments used games with monetary payoffs presented to the participants and decision-making was not movement-related. Some tasks were presented with a spatial dimension (e.g., Isoni et al., 2013, 2019), but relative locations were used as payoff-irrelevant labels, rather than parts of the payoff structure. Here, we treat spatial dimensions as payoff-related in a bid to extend the investigation of focal points to the action domain. After summarizing relevant findings from the behavioral economics literature, we present the hypothesis of the present study, outlining briefly the experiments that tested our hypothesis.

5.1.1 Focal points

In the context of mixed-motive games, Schelling (1960) analyzed situations of tacit bargaining (i.e., coordination of decisions without explicit communication) between interaction partners with varying degrees of common interest. As he pointed out, even in war and business, opponents always have *some* level of shared interest, since they both aim to reach an agreement, so coordinating on any profile of strategies is better than not coordinating at all (Schelling, 1960). In

[31] This is not a new idea. Lewis (1969/2008) also acknowledges that there are variations of the meeting-place problem that are not pure coordination problems but mixed games. Interaction partners might not derive the same utility from all potential meeting points. Some locations may be more "valuable" than others. This idea is also reflected in Schelling's (1960) parachutist example (pp. 58-59).

[32] N.B. A few exceptions are research in the motor domain that investigated coordinating on Nash equilibria in continuous movement task versions of games from classical game theory (Braun et al. 2009, 2011).

tacit coordination, people mutually predict the expectations of others about their future actions and match these expectations. However, predicting expectations could devolve into an infinite regress[33], and negotiation must stop at one point: expectations must converge. A focal point is a combination of strategies that provides a unique solution among equivalent or similar alternative solutions in a game, in other words, it enables the selection of a Nash equilibrium[34,35] in a game with multiple equilibria (Bacharach, 2006; Schelling, 1960).

5.1.1.1 Pure coordination games and label-salient focal points. According to Schelling's (1960) focal point theory, salient non-payoff related characteristics are unique labels attached to strategies that are "prominent" (Schelling, 1960; pp. 57-58, 64, 69) to the decision-maker because of associations or analogies that others will likely also be aware of. For example, two parachutists might find each other after landing separately by converging on the most conspicuous landmark on a map they both possess (Schelling, 1960, pp. 54-55). Making strategic decisions based on payoff-irrelevant labels (*label salience*, Crawford et al., 2008) goes against classical game theoretical assumptions that derive rational decisions on purely mathematical bases (Nash, 1953) – but in symmetric games like the "heads and tails" matching game, classical game theory is unable to predict the overwhelming coordination on "heads" (Bacharach, 2006). This discrepancy intuitively suggests that there is something beyond payoff structure that can help coordination when communication is impossible (Bacharach, 2006; Schelling, 1960).

Indeed, in pure coordination games where participants have to choose matching strategies and neither player can be disadvantaged in terms of payoffs based on their chosen strategies in any of the possible equilibria, participants were found to use salient label cues for coordination. They

[33] E.g., person A predicts that person B expects her to expect B to expect her to choose X and so on. Level-k reasoning and cognitive hierarchy theory are alternative approaches to decision-making in coordination problems, which rely on different reasoning stages and participants making assumptions about the level their partner might reason at (Camerer et al., 2004; Stahl & Wilson, 1995). The probability of using focal point-reasoning versus level-k/cognitive hierarchy reasoning may depend, among others, on the type of game played (Bardsley et al., 2010; Faillo et al., 2017).

[34] A Nash equilibrium is a combination of strategies from which neither player can deviate unilaterally without suffering a loss in payoffs (Nash, 1951).

35 More recent research found that non-equilibrium strategies may also become focal points when their relatively lower payoffs are unique among multiple identical Nash equilibria (Bosch-Domènech & Vriend, 2013).

successfully coordinated their decisions with higher probability than random behavior would predict (Bacharach, 2006; Crawford et al., 2008; Mehta et al, 1994a, 1994b). Salient label cues can be color, closeness of concepts, aesthetic principles like balance in an image (Mehta et al., 1994a, 1994b) or well-known landmarks (Crawford et al., 2008). The notion of salience refers to properties of strategies that are unique "in some conspicuous respect" (Lewis, 1969/2008, p. 35) and can be characterized as one of multiple types: primary, secondary, or Schelling salience (Mehta et al., 1994a; 1994b; based on Lewis, 1969/2008).

A label has primary salience if it comes to mind for a player when faced with the coordination game, possibly due to a stochastic process and not due to reasoning (Mehta et al., 1994a). As such, primary salience is part of non-rational play[36]. Secondary salience, on the other hand, can be considered bounded rational (Mehta et al., 1994a). When a player chooses according to secondary salience, she anticipates which label might have primary salience to her partner in the game and try to maximize her own utility by choosing that label[37]. Secondary salience is expected to lead to more successful coordination than primary salience, although the ranking of chosen labels should be similar based on the two types.

Schelling salience is different from the first two types (Mehta et al., 1994a; Schelling, 1960). Using Schelling salience refers to coordinating decisions based on a selection rule which can lead to successful coordination if both players recognize it and apply it to the problem (Mehta et al., 1994a). For example, when participants are required to pick any number for coordination, they tend to converge on 1, which is the smallest positive integer. Presumably, this description is used as a selection rule as it unambiguously identifies a solution from among an infinite amount of numbers (Schelling, 1960). Schelling salience may coincide with secondary salience (e.g., when the rule is "Choose the label that is likely to be chosen by others") but also inspires different predictions

[36] However, if the same label has primary salience for a sufficient number of people in a population because of their common cultural background or experiences, primary salience can lead to more coordinated decisions than would be expected based on random choices.
[37] This salience type can be extended to *n*-th level salience (where a player anticipates which strategy would have secondary, tertiary etc. salience to her partner, Mehta et al., 1994a).

than primary and secondary salience (Bardsley et al., 2010; Mehta et al., 1994a). Mehta and colleagues (1994a) found that people are able to coordinate their decisions at a level higher than a primary salience-based decision strategy would suggest. Although they did not disambiguate secondary and Schelling salience in their experiment, they argued that Schelling salience likely plays an important role in solving coordination games.

5.1.1.2 Payoff-salient focal points and team reasoning. The focal points described above are independent of the mathematical structure of a game (Schelling, 1960). However, people may also converge on a payoff-dominant equilibrium that ensures the highest payoffs for both players (Harsanyi & Selten, 1988), that is, people can use *payoff salience* (Crawford et al., 2008) for coordination. Payoff-salient focal points and tacit bargaining ("virtual bargaining") were proposed to play important roles also in social decision-making and joint actions (Misyak & Chater, 2014; Misyak et al., 2014).

The Hi-Lo game analyzed by Bacharach (2006) is a good example of joint payoffs acting as focal points. In this game, the players have to match their choices with one another by choosing "high-high" or "low-low". If both choose "high", they each get the same amount of rewards, which are larger than the amounts they would receive if they both picked "low" (i.e. high-high is payoff-dominant). In case they choose non-matching labels, their payoffs are zero. Bacharach (2006) addressed the paradox that the "high-high" solution appears intuitively *right*, even though classical game theory does not predict that it is better than "low-low". He argues that players come to the "high-high" solution by adopting a distinct form of reasoning: team reasoning.

According to Bacharach (2006), team reasoning is deployed by individuals that identify with a group and select solutions in games that will maximize the payoffs of the group, rather than of the individual. Having identified the joint-payoff maximizing equilibrium, the decision-making player executes her part in that strategy profile (Bacharach, 2006). A similar theory of solving coordination problems as a team was also formulated by Sugden (1993), who emphasized that players look for the *best rule* to coordinate on a solution that will be recognizable for everyone and lead to coordination. Some experiments suggest that people engage in team reasoning and select joint utility-maximizing solutions in coordination games (Colman et al., 2008a), and can coordinate on payoff-dominant strategy combinations that ensure the highest payoffs for both players (Bacharach, 2006; Bardsley et al., 2010).

It is worth noting, however, that other experiments addressing coordination on payoff-dominant strategies found that people did not consistently coordinate on equilibria that yielded the highest payoffs for the players (Cooper et al., 1990; Van Huyck et al., 1990). Participants coordinated either on payoff-dominant equilibria, on cooperative strategies that were not equilibria[38] but ensured high payoffs for both players (Cooper et al., 1990), or failed to coordinate on any equilibria initially and later coordinated on low-risk, low-payoff equilibria (Van Huyck et al., 1990). It seems that there is no unequivocal evidence for people coordinating on strategy pairs that could serve as payoff-maximizing focal points.

5.1.1.3 The influence of payoff asymmetries on the power of focal points. So far, we focused on coordination games without conflicts of interest between the players. What happens when a strategy pair with asymmetric payoffs could serve as a focal point? Schelling argued for the validity of label-salient focal points in bargaining games where equilibria resulted in conflicting payoffs for interaction partners, that is, in games that are not pure coordination games. However, some experimental evidence suggest that uneven distributions of payoffs in focal equilibria reduce people's coordination based on focal point-reasoning, both when strategy pairs were focal because of salient label cues or because of payoff salience (Crawford et al., 2008; Faillo et al., 2017; Isoni, et al., 2019; Parravano & Poulsen, 2015; Poulsen & Sonntag, 2019; van Elten & Penczynski, 2020; but see some mixed results from Isoni et al., 2013; 2018). Interestingly, people also relied less on salient label cues when exact payoff-related information was not provided in the game and there was uncertainty about the asymmetries between participants – where an increase of label-based focal point use was expected (Isoni et al., 2019).

[38] I.e., a player could deviate from the strategy pair for a better payoff, given the other player's choice.

Recently, Isoni et al. (2018) suggested that the negative effect of payoff asymmetry was due to conflicts of interest between people that caused the players to prefer different rankings of strategy pairs, rather than the payoff asymmetry itself. When the asymmetry in individual payoffs is present in a coordination game but the conflict of interest in strategy ranking is eliminated, payoff inequality does not hinder coordination (Isoni et al., 2018). It might even increase coordination success: Gueye et al. (2020) found that a motivation to gain the highest total payoff was behind the coordination success in a high payoff asymmetry condition. In summary, the literature appears inconclusive on whether and how focal points benefit coordination in economic games if there are asymmetries between the individual payoffs of interaction partners. Some evidence points to the power of joint payoffs as focal points even if the players gain unequal payoffs, others suggest that payoff asymmetries reduce coordination on focal points. Furthermore, to the best of our knowledge, how this might play out in the action domain remains unexplored.

5.1.2 Research question and hypothesis

Once the NYC meetup game becomes a practical problem subject to execution, it might be better characterized as a tacit bargaining game similar to the parachutists' problem (Schelling, 1960) due to the expected costs of movement in the action planning process, and the potential asymmetries between them. Schelling argued (1960, pp. 58-59) that conspicuous landmarks still serve as the best coordination devices for parachutists who dislike walking. We test an alternative solution by addressing the question whether joint-payoff related focal points help coordination in joint actions under uncertainty, in spite of cost asymmetries between the co-actors[39].

In our previous studies reported in Chapters 2-4, we established that utility-maximizing decision-making is relevant to sequential joint action contexts where participants have to plan movements by taking into account implicit, movement-related costs. Building on those results, we test the hypothesis that, beyond its instrumental benefit to the execution of joint actions, co-

[39] We refer to cost asymmetries in what follows, since payoffs were primarily manipulated through the distance-based costs of movements. In addition, successful coordination would be "rewarded" with 1 point for each actor, whereas miscoordination would result in 0 points (see section 5.2.5. Procedure on the experimental procedure).

efficient decision-making provides focal points to interaction partners in situations of uncertainty about a co-actor's decisions.

We propose that, on the one hand, co-efficiency could act as a Schelling salient rule for selecting an action plan that increases the probability of successful coordination. The rule could be stated as "Choose the object reachable by the action option with the lowest expected joint cost of movements". Such a group-focused rule for decision-making would provide more useful cues for coordination than focusing on individual-level movement costs. This is because aiming to maximize the joint utility of a group restricts potential coordination points to a greater extent than appealing to individual payoffs of any agent (Bacharach, 2006). Furthermore, in problems where the potential points for coordination have no other salient characteristics than visual object features (e.g., shape) and the movement costs related to reaching them, using the joint movement costs as focal points can sometimes be more discriminative than using object features[40]. The joint-cost minimizing feature, on the other hand, would also make co-efficiency a payoff-salient focal point. Therefore, according to our hypothesis, joint action costs would be both payoff-salient and Schelling salient, creating a link between Schelling's idea of label-salient focal points and research on coordination on payoff-dominant solutions.

We predict that if co-efficiency is a potential focal point that people might use for coordination, then in a remote coordination version of the object matching task from Chapter 3, we will find that the participants' choices are more often consistent with a co-efficiency maximizing (i.e., joint-cost minimizing) decision strategy than random choices would predict. Further, this tendency could be as strong as, if not stronger than, a strategy based on object shapes or other non-payoff-salient coordination strategies.

[40] Take, for example, the situation where two parachutists could meet either by a certain tall tree, a small tree, or a thick tree, each in a separate area of a park. The probability of coordinating based on the shape of a tree is 1/3, whereas if the parachutists applied the rule of minimizing the total amount of walking their dyad would have to undertake to reach a meeting point, they could easily infer the tree by which both of them would meet, with probability > 1/3, provided that the joint costs of walking would not be equal between the three trees' cases.

5.1.2.1 The experiments. We present two online experiments addressing these hypotheses. In both experiments, the participants were told that they were taking part in an object matching game with a remote partner, with whom they could not communicate (in fact, the participants took part in the study on their own and there was no remote partner). We used stimuli generated by the code of the task used in Chapter 3 to instantiate a modified online version of the object matching task. The participants were informed that they were provided the same information about the task as the remote partner. They were instructed to choose objects in a virtual environment that they thought would match the remote partner's choice. No feedback was given to them about whether they successfully picked the same-shaped object as their partner. Each trial differed in terms of the individual and joint action costs the participants would have to incur when moving with their mouse to collect an object by clicking on them.

The two experiments differed only with regards to the method of object collection, which placed differential emphasis on the joint movement costs in a trial. In Experiment 1, the participants had to complete a three-step movement sequence to collect an object, which was closely based on the task in our previous lab-based Study 2 (Chapter 3). This procedure was implemented to provide sufficient information for the calculation of the expected individual and joint costs of the onscreen movements that actors had to perform. This experiment tested the hypothesis that people would try to coordinate by choosing objects that minimize the joint costs of object collection.

Experiment 2 tested the hypothesis that the effect we found in Experiment 1 was due to the representation of joint movement costs, and not merely to the visual features of the onscreen arrangements of objects. We instructed the participants to collect an object by simply clicking on them once. This one-step procedure was implemented to make the calculation of the expected individual and joint movement costs as defined in Experiment 1 unviable. Every object collection action could now start from any onscreen location of the cursor, and by moving directly to the chosen object and knowing that the remote partner was doing the same, calculating the joint movement costs was no longer possible in Experiment 2. An effect of co-efficiency on people's decision-making would suggest that the visual configuration served as focal point, rather than the action costs associated with the task.

5.2 Experiment 1 - Methods

5.2.1 Participants

We recruited participants through Prolific, a recruitment service for online studies in the social sciences (www.prolific.co). The participants gave their informed consent and received monetary compensation through Prolific in exchange for their participation. The study was approved by the United Ethical Review Committee for Research in Psychology (EPKEB) in Hungary. Participants were screened according to the following criteria: they were right-handed, native English speakers between the ages of 18-100 years, with normal or corrected-to-normal vision. All of them had previously indicated that they would be comfortable with participating in studies that employ deception, with the condition that they would be debriefed about the true nature of the research post-experiment.

In our first study addressing co-efficiency (Chapter 2), we analyzed the decision data of 12 dyads per experiment using Wilcoxon signed-rank tests to compare choice proportions to chance. We had no prior expectation of the size of the co-efficiency effect we could find in the novel online setting. Therefore, to ensure that the present study was adequately powered, we aimed, at a minimum, to replicate one of the smaller effects we found previously. A power analysis in G*Power 3.1 (Faul, Erdfelder, Lang, & Buchner, 2007) estimated that a sample size of 26 would be necessary to provide 95% statistical power to achieve the same effect size (Cohen's $d = 0.810$). Due to the convenience of online participant recruitment, we opted to test a sample of 50 participants, exceeding the minimal necessary sample size of 26.

In total, 68 participants took part in Experiment 1. We analyzed the data of 50 participants (23 males, age $M = 34.9$ years, $SD = 12$). Criteria for data exclusion were as follows. We would exclude and replace the data of participants who did not pass any of three requirements: (1) answering correctly on at least 75% of the attention check trials (= 3 trials), or (2) completing at least 75% of the experimental trials (= 24 trials) by making legitim object choices, i.e., clicking on one of the objects available to them in a trial, or (3) completing at least 75% of the task by adhering to the instructions for object collection. Participants who complied with at least the first two criteria were awarded a fixed bonus of 2 GBP on top of their base remuneration for their time (3.75 GBP/30 minutes allocated to the task). We excluded and replaced the data of 18 participants who did not complete at least 75% of the task by adhering to the instructions for object collection[41], even though they answered correctly in the required minimum amount of attention check trials and made legitim object choices in at least 75% of the experimental trials.

5.2.2 Apparatus

The experiment was built and hosted on Gorilla (www.gorilla.sc), a website for creating online behavioral and survey studies. Each participant used their own computer and internet browser in the convenience of their own home. The use of tablets and smartphones was excluded on the study's recruitment page, so that participants used a laptop or a desktop computer. Participants were asked to open the study in Google Chrome or Firefox, and to use a mouse instead of a touchpad, if possible. These requests were made to reduce the variety of software and hardware that participants used, and to limit the range of different screen sizes.

The majority of participants used a computer mouse (n = 36), some of them reported using a touchpad (n = 14). Gorilla collected information about the size of the active screen area a participant used to complete the experiment. Active screen area sizes ranged from 1263 x 578 pixels to 2560 x 1255 pixels.

[41] These participants made their choices either by (1) clicking three times on the instructed objects in the wrong order (n = 6), or (2) clicking only on one kind of object multiple times (n = 3), or (3) clicking only twice, on a combination of the specified objects (n = 7), or (4) clicking only once, on one of the target objects (n = 2).

5.2.3 Stimuli and Task

The present study used a list of spatial layouts showing a simple virtual environment. In each trial, a 750 x 750 px image of a layout with the following objects was displayed to the participant in the middle of her browser window: (1) a thin black wall dividing the screen into two halves, corresponding to the two actors' task areas, (2) two pairs of black target objects (two circles and two squares, 21 x 21 px) distributed between the two task areas (one of each shape displayed on each side), and (3) two black-bordered octagonal starting locations (67 x 67 px) with another, smaller octagon inside (42 x 42 px, Figure 5.1). The starting locations were always displayed at mid-position (375 px) along the horizontal sides of the layout.

Gorilla automatically adjusted the size of the layout images to the size of the participant's active screen area: if the active screen was smaller than 750 px in height, images were downscaled, for larger screens, the images were kept at the original size. Importantly, aspect ratio was always preserved. Therefore, the distance relations between each black target object and the octagons were constant, regardless of absolute active screen size.

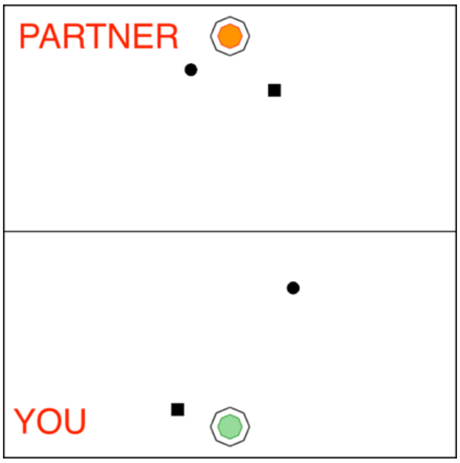

Figure 5.1: An example of the layouts in the online object matching task. Starting locations were indicated by the two black-bordered octagons and the smaller colored octagons within, and the locations for the two pairs of black target objects (circles and squares) were generated by stochastic selection processes. The red labels in the image were only visible in the example layouts shown in the task instructions, not in the task.

Throughout each trial, the color of the smaller octagon by the participant's side of the layout changed. At first, both actors' octagons were orange-colored to signal a decision-making phase (Figure 5.3 illustrates the temporal structure of the trials). To cue the participant to collect the chosen black target object, her octagon by the lower side of the layout changed to green, whilst the remote partner's octagon remained orange.

To filter out participants who did not pay sufficient attention during the task, four attention check trials were included in the experiment. These were two alternative forced-choice tasks where participants had to identify animals in photographs, one in each trial. Participants had to choose between two labels and could answer in their own time – e.g., they saw the image of a parrot and had to click on either the label "Parrot" or "Chicken".

5.2.4 Design

The central assumption of the design is that the cost of an action is a monotonic function of the path length that has to be covered onscreen by the participant's hand (i.e., the cursor) to collect an object in the task. Three cost disparity parameters were manipulated to test the influence of individual versus joint action costs on decision-making. The Joint cost disparity parameter quantifies the difference in the summed movement costs of a dyad between the two matching object pairs available to them in an experimental trial (see Figure 5.1). This parameter is the sum of the individual Self and Other cost disparities in the trial, each of which quantify the difference between the distances of a participant's (and her partner's, in case of the Other cost disparity) starting octagon to each of the black objects on the participant's (and her partner's) side of the screen.

The disparity parameters were calculated by subtracting distances to black circles ("b1" for the participant and "b2" for the partner) from distances to black squares ("a1" for the participant, "a2" for the partner), both on the participant's and the partner's screen side. Self Disparity was calculated as a1-b1, Other Disparity as a2-b2, and Joint Disparity as the sum of the two: (a1+a2)-(b1+b2) = (a1-b1)+(a2-b2). The decision to subtract from the distances to square objects is arbitrary but important to keep in mind as the dependent variable was the probability of the participant choosing the black square (object "A1") in their side of the layout.

In the present study, the distributions of the Self and Other cost disparities were statistically independent from each other across trials (Figure 5.2a). We sampled Self Disparity and Other Disparity for each trial independently from the same uniform distribution (between -265 and 265 pixels). As a consequence, Joint Disparity had a triangular distribution and was positively correlated with both terms (Figure 5.2b-c).

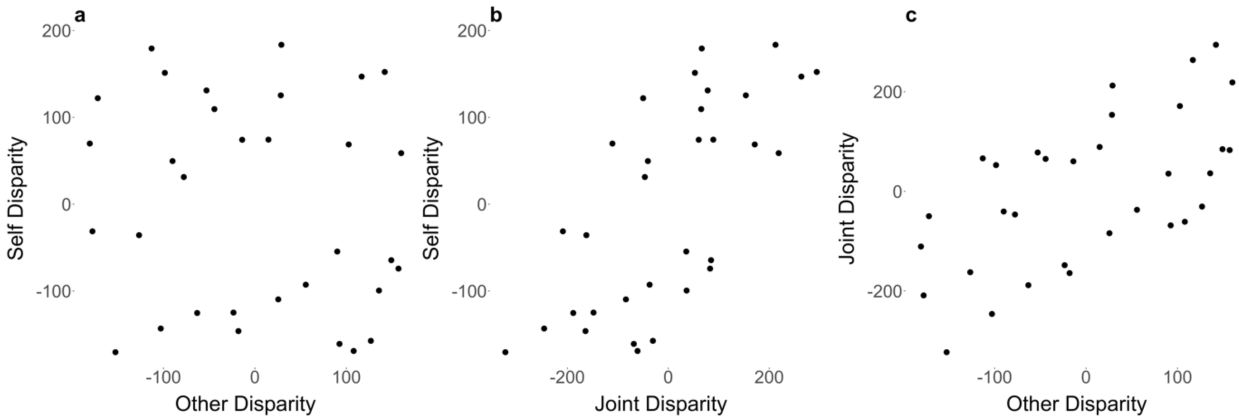

Figure 5.2: Scatterplots of the joint distributions of cost disparities used in Study 4. (a) Self and Other Disparities were uncorrelated, $r = -.146$. (b) Self and Joint Disparities, and (c) Joint and Other Disparities were positively correlated with each other ($r = .691$ and $.614$, respectively).

We generated a list of 32 trials with the additional constraints that the (1) black square in the participant's task area appear on the left and right side an equal amount of times (16-16 trials), and (2) the square would be the joint-cost minimizing solution in 16 trials, implying that in 16 trials, the circle pair would minimize the joint action costs. A third additional constraint related to the congruency between individually and jointly (co-)efficient action plans in a trial.

The de-correlation of Self and Other cost disparities resulted in an imbalanced generation of congruent and incongruent trials. Due to the positive correlation between Self and Joint disparities, the majority of trials generated would be classed as congruent: the individually efficient, path-minimizing, choice for the decision-maker would also be efficient for the group (joint-cost minimizing). A large subset of these trials was "joint-other" congruent, i.e., optimal predictions for the group were not only optimal for the decision-maker but also for their partner, both Self and Other-cost minimizing (see details in Appendix B, B.1.1 Additional information on congruency (section 5.2.4 Design)). To minimize the sampling of congruent trials as much as possible, we added the third constraint of generating congruent trials 75% (= 24 trials) of times in the whole experiment. Of these 24 congruent trials 14 trials were joint-other congruent. Therefore, participants saw 10 joint congruent, 14 joint-other congruent, and 8 incongruent trials.

5.2.4.1 Side imbalance. Due to the random generation of the trials and because we did not specifically constrain this, there was an imbalance in the occurrence of the object in the co-efficient position on each side of the layouts. In 18 trials out of 32, the co-efficient object was displayed on the left side (9 squares, 9 circles), and on the right side in 14 trials (7 squares, 7 circles).

5.2.4.2 Trial presentation. The list of trials was presented to the participants in a fully randomized order, divided into 4 blocks of 8 trials each. Among the experimental trials, every block featured one attention check trial (animal image identifications) in the first half, the middle, or the end of the block. Participants could take breaks between blocks, however, once a block was started, the experimental trials proceeded independent of the participant's actions within blocks.

5.2.5. Procedure

After they decided to take part in the study on Prolific, participants were redirected to Gorilla. They gave their consent to participation and read the instructions at their own pace by clicking through a series of pages. They were informed that they would be connected to another participant online, with whom they would play an object matching game in real time. This was part of an experimental deception, as participants completed the task on their own.

The two alleged participants could play one of two roles: South player (Actor 1) by the lower half of the screen and North player (Actor 2) by the upper half of the screen. Participants were always assigned to play the role of the South player. They were instructed to choose the object in each trial that they think would match the shape chosen by the remote partner, without seeing the partner's actions or being given any feedback about their choices. It was made clear to them that the North player did not have information about the participant's actions either but had access to the same information in the task as the participant – i.e., the co-actors saw the same layout in each trial and were given the same task instructions. The instructions stressed the importance of coordinating with the partner's choice by stating that each 'matching' object choice would earn 1 point to both players. This was also part of the deception, designed to generate a joint goal and to motivate a cooperative mindset in the participants.

After reading the instructions, the participants completed four practice trials before proceeding to the task. At the beginning of each trial, first, a fixation cross was displayed on the screen for 1500 ms, before and after 100 ms of a blank screen display. Then, participants were shown the layout with the objects and two orange-colored octagons for 3000 ms (Figure 5.3 – Decision phase). During this phase, the smaller octagons were both orange-colored to signal to the participant that she was required to decide which black object she would collect, a square or a circle on her side of the layout (lower half).

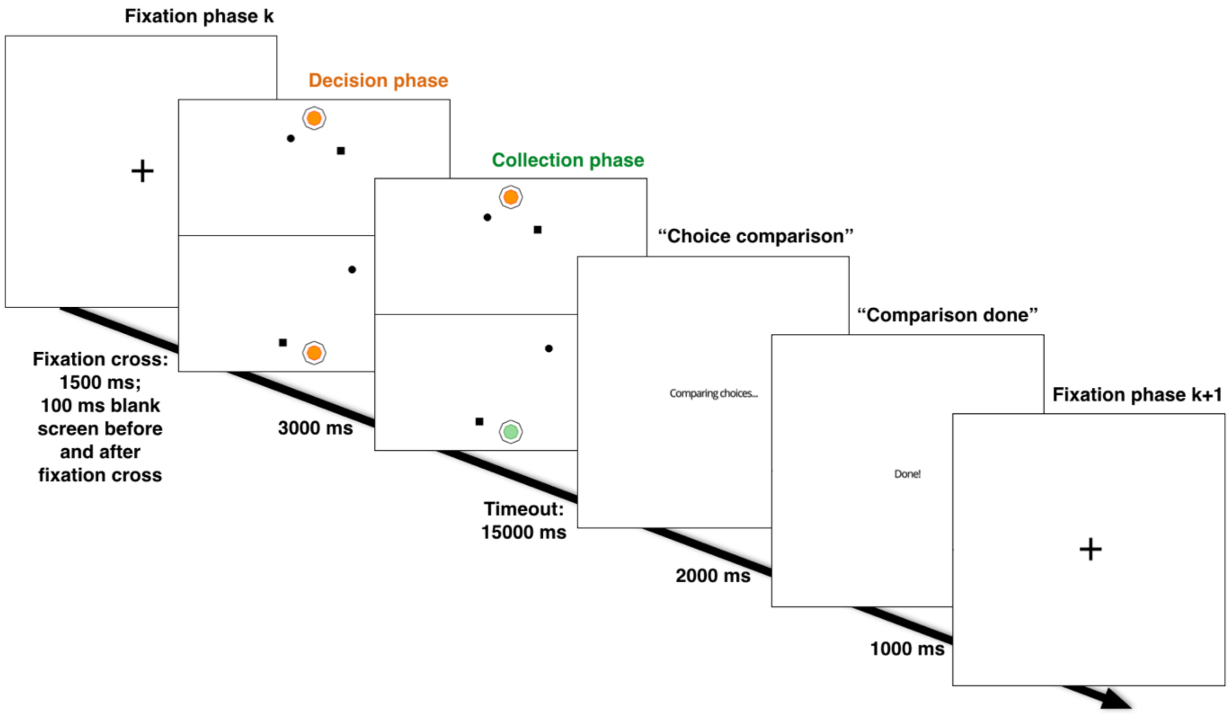

Figure 5.3: The temporal structure of an experimental trial in Experiment 1 and 2.

After 3000 ms, the color of the octagon by the lower border of the layout switched to green, which served as a cue for the participant to start collecting the chosen object (Collection phase). Collection of the objects consisted of the participant clicking on the three relevant regions of the screen in order: green octagon – the chosen black object – green octagon. A 3 px red dot appeared on every click as an action feedback. The red dots helped participants track the number and location of their own clicks. This phase of the trial timed out after 15000 ms, after which time the next phase of the trial was automatically displayed. Once participants collected a black object in the Collection phase, a new message informed them (displayed for 2000 ms) that their choice was being compared to their partner's decision. This was followed by the text "Done!" for 1000 ms. Afterwards, a new trial started with a fixation cross.

Following the behavioral task, the participants filled in a questionnaire on their experience playing the object matching game. They provided demographic data and answered questions on a Likert scale (1 – "Not at all" or "Not at all likely"; 7 – "Very much" or "Very likely") regarding their perceived coordination difficulty and success, as well as how successful they thought their partner was in matching their choices with the participant's. We questioned them on any conscious coordination strategies they may have used, as well as their belief in the presence of a remote human partner. They were debriefed about the deception in the study and directed back to Prolific to submit their data to the researcher. The participants spent on average M = 16.4 minutes (SD = 5.77) on completing the task and filling in the questionnaire.

5.2.6 Data Analysis

The primary dependent variable was the object choice made by the participant. Gorilla registered the coordinates of the participant's clicks in the Collection phase, which enabled the offline coding of object choices.

5.2.6.1 Choice proportion analyses. Object choices were coded in terms of (1) shape (square or circle), (2) position in the layout (left or right side of the screen), (3) relative distance from the participant's octagon in pixels (short or long individual path taken to an object), (4) distance from the cursor's position at the beginning of the trial's decision phase, and (5) co-efficiency (co-efficient or sub-efficient option for the virtual dyad). The proportions of choices along the first four dimensions were tested against chance level (.5), using two-tailed Wilcoxon signed-rank tests. The proportions of co-efficient object choices were compared to a chance level of .75, in addition to chance level at .5, due to 75% of the trials being congruent (see section 5.2.4

Design2), where joint-cost minimization coincided with self-cost minimization. We also compared to chance (.5) a measure of expected coordination success between the participants in the sample. We report results of correlational analyses (Spearman's ϱ) to describe the relationships between the aforementioned object choice types and coordination success. We used Bonferroni-corrections to adjust for multiple comparisons[42]. *V* statistics are reported for the Wilcoxon tests, as well as matched-pairs rank-biserial correlation coefficients (*r*) as measures of effect size (Kerby, 2014).

5.2.6.2 Logistic regression models and model comparison. We used the cost disparity parameters (in pixels) in each trial to predict the probability of the participants choosing the square ("A1") using generalized linear mixed models (GLMM) fitted with the *glmer* function from the *lme4* R package (Bates et al., 2015). In every model, the likelihood of the participants choosing A1 was predicted by different cost predictors and their combinations through a logistic link function. We predicted negative β coefficients for all cost parameters that significantly affected the predicted probability. This statistical approach provided group- and individual-level estimates for the parameters included in the mixed models. We tested (1) the co-efficiency as focal point hypothesis and (2) alternative hypotheses that participants may have used coordination strategies centered on selfish or altruistic choices (i.e., Self- or Other-cost minimizing strategies), or not trying to coordinate (i.e., picking the object closer to the cursor's location in the decision-making phase).

We estimated five primary mixed-effects logistic regression models: (1) *Self Disparity*, (2) *Other Disparity*, (3) *Distance from Cursor*, (4) the linear combination of *Self and Other Disparities* reflecting a joint-cost minimization strategy, and (5) the linear combination of *Self and Other Disparities and Distance from Cursor*. *Subject ID* was included in every model as random effect grouping factor to account for individual differences and dependencies among data points within individuals (Singmann & Kellen, 2019). By-subject random intercepts and random slopes were estimated.

[42] The α-levels for Wilcoxon signed-rank tests were divided by seven (.007): tests were run on the four potential biases mentioned in the text, on the ratios of co-efficient choices against 50% and 75%, and on the coordination success measure. The α-levels for Spearman's ϱ tests were divided by five (.01): we ran these to examine the relationships between expected coordination success and the co-efficient choices and the four other potential biases.

Additional models. We also fitted extended versions of each model where we included the factor *Side*: the side of the screen on which the square object was located in a trial (left v. right). We fitted these models to examine whether a weak tendency to pick the object located on the left side was in itself a coordination strategy or whether it was a by-product of making co-efficient choices (see section 5.2.4 Design1 on Side imbalance). The results of these estimations are available in Appendix B, as Side was not found to exert significant influence on decision-making in models with high predictive accuracy (Table B.7-12).

We also estimated models including the following predictors: (1) Square categorized as Co-efficient vs. Sub-efficient object (Yes/No) (Table B.6), the linear combinations of (2) Self Disparity and Distance from Cursor, and of (3) Other Disparity and Distance from Cursor (Table B.13-14), and (4) Intercept only (Table B.15). In total, we fit 15 regressions.

Model comparison. We aimed to select the model that best describes the observed data and is expected have the highest accuracy at predicting future behavior. These two goals trade off against each other, since more complex models with more parameters usually better fit the data that they were estimated on, but complexity could result in overfitting which in turn can be disadvantageous for predicting data on future samples (McElreath, 2020). We ranked models based on metrics separately for the fit of models to the data (R^2), and for predictive accuracy (AIC values and derived Akaike weights). We report both marginal[43] and conditional R^2 values (Nakagawa & Schielzeth, 2013) in Appendix B but focus only on the marginal R^2 for model comparison here as our main interest lies in the fixed effects of the cost parameters.

We selected the model that was expected to provide the highest accuracy in predicting future data based on their Akaike Information Criterion (AIC; Akaike, 1973) and Akaike weights[44]. The

[43] The marginal R^2 quantifies the proportion of variance explained only by the fixed effects, whereas the conditional R^2 is computed based on the combination of fixed and random effects (Nakagawa & Schielzeth, 2013).

[44] The Akaike weights are the normalized evidence ratios between each model *i* and the model with the smallest AIC. The weights express the conditional probability that model *i* is the best model that minimizes expected K-L divergence given the observed data and the set of candidate models (Burnham & Anderson, 2004; McElreath, 2020; Wagenmakers & Farrell, 2004). Note that the weights we report were calculated based on all 15 models, including those reported only in Appendix B.

AIC quantifies the amount of expected deviance of the model from potential new data (i.e., an estimate of the distance between the true state of the world and the model approximating it, Burnham & Anderson, 2004), and it is calculated by correcting the estimated maximum log-likelihood of a model with a bias correction term proportional to the number of parameters in the model to penalize for complexity (Burnham & Anderson, 2004; McElreath, 2020). The lower the AIC, the better a model's predictive power relative to other candidate models. We compared the Akaike weights (w_i(AIC)) to identify the relative best model given the observed data (Burnham & Anderson, 2004; Wagenmakers & Farrell, 2004).

5.3 Results

5.3.1 Descriptive Statistics

5.3.1.1 Coordination Success. Previous studies calculated a mean expected coordination success metric either using a legitimate matching process (Isoni et al., 2013, 2019), or by calculating a coordination index for each question collapsed over participants (Mehta et al., 1994a). In the present study, coordination success could not be determined as actual matches between objects chosen on the two sides of the screen, since all participants were tested on the same layouts in the same orientations. We calculated a measure of mean *expected coordination success* by pairing each participant with all the remaining participants in the sample (n = 49) and checking whether on each trial, the pair chose the same objects. We calculated the ratio of success for each virtual pair (e.g., 17/32 for 17 matches out of all trials), and for each participant we averaged the ratios of expected matches with others to yield participant-level mean expected coordination success measures. These were compared with a Wilcoxon signed-rank test to chance. Chance level performance would suggest that participants were not converging on any coordination strategy.

On the group level, participants' mean expected coordination success with one another was statistically significantly higher than 50%, and this appeared to be a relatively large effect (Figure 5.4a, M = .56, SD = 0.06, V = 1095, p < .001, r = .72, 95% confidence interval (CI) for the median choice proportion .56 = [.54, .60]).

5.3.1.2 Co-efficient choices. According to our main hypothesis, choosing the objects in the co-efficient position could serve as a useful coordination strategy in the present task. We found that overall, participants chose the object that was in the co-efficient position in 998 out of 1585 trials[45], and the ratio of co-efficient choices was significantly higher than 50% (Figure 5.4b, M = .63, SD = 0.14; V = 902.5, p < .001, r = -.42, 95% CI for proportion .66 = [.60, .72]), suggesting that overall, there was a moderate tendency to choose the object reachable by the co-efficient total path over the alternative. However, we set chance level at 75% to account for the possibility that participants could have made co-efficient choices for Self- or Other-cost minimizing reasons in the congruent trials. The ratio of co-efficient choices was statistically significantly lower than .75 (V = 75, p < .001, r = -.88, 95% CI for proportion .62 = [.56, .67]), suggesting that co-efficient decisions may have been made for selfish or altruistic (Self- or Other-cost minimizing) reasons. The regression models provide a more detailed report on this. We found that the proportion of co-efficient choices strongly correlated with the expected coordination success of the participants (Figure 5.4c, Spearman's ϱ = .746, p < .001), suggesting that the more often people opted for the object in the co-efficient position, the likelier it was that their choices matched with other participants 'choices.

[45] The total number of trials in Experiment 1 was 1600, however, the proportions of choices were calculated based on trials with no missing values (i.e., where a valid decision was made). Missing values were defined as trials where participants clicked anywhere other than the two black objects available to them in their part of the screen. In total, the data had 15 missing values.

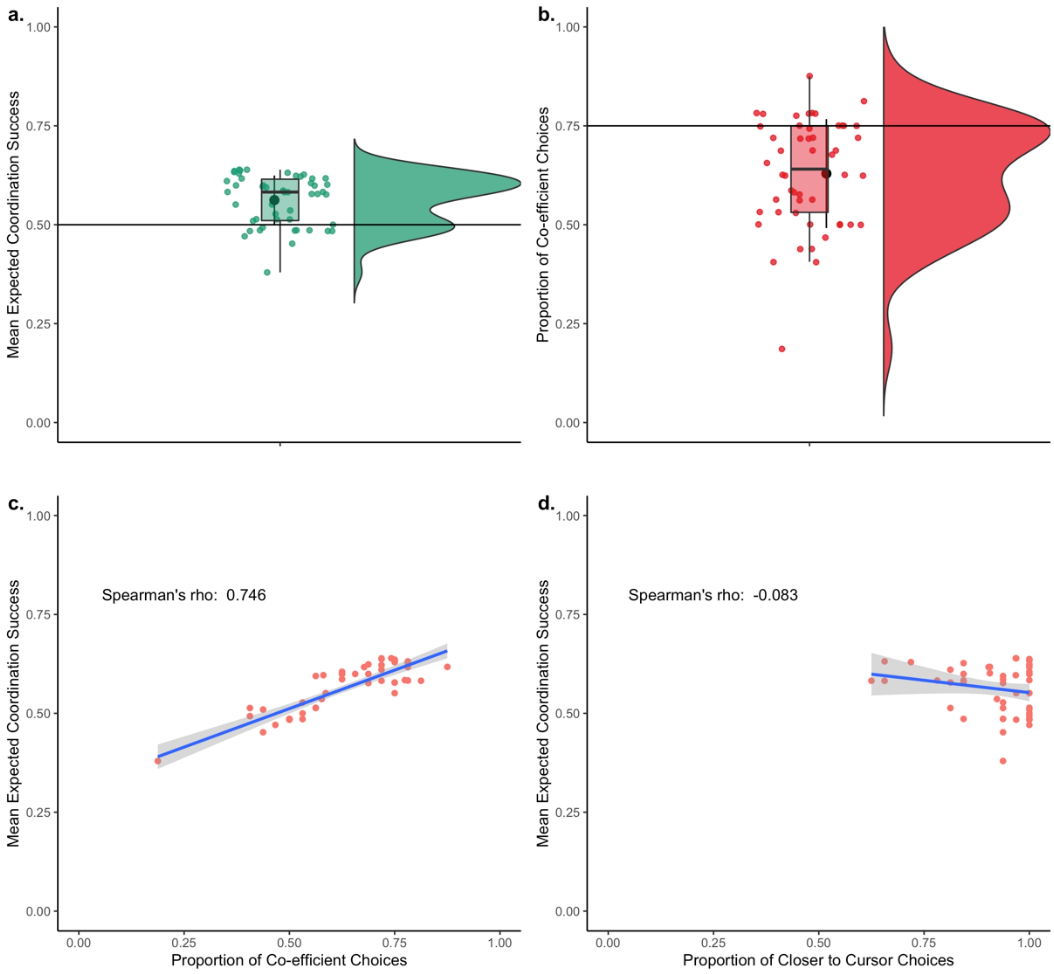

Figure 5.4: (a) Mean expected coordination success and (b) proportions of co-efficient object choices in Experiment 1. In each boxplot, black horizontal lines indicate medians, the bar's whiskers extend 1.5 times the interquartile range; black dots signal the means, with their whiskers extending to one SD; and the density plots illustrate the distributions of (a) mean coordination success and (b) choice proportion data. Each colored dot represents a participant's choice ratio according to the given variable. The horizontal line indicates chance level (.5 for coordination success and .75 for co-efficiency). Correlations between mean expected coordination success and (c) the ratio of co-efficient choices, and (d) the ratios of choosing the object closer to the mouse cursor at the beginning of a trial. The bold blue lines are fitted linear model predictions, the shaded area is a 95% confidence interval.

5.3.1.3 Object Shape, Side, Short Path Biases. To investigate whether participants used any alternative strategies in the task, we compared the proportions of choices along dimensions other than co-efficiency to chance (.5), with α-levels set at .007.

'Object shape' bias. The participants chose object A1 (the black square on their screen side) 794 times out of a total of 1585 trials. The proportion of square choices was not statistically significantly different from chance (Figure 5.5a, M = .50, SD = 0.22; V = 426, p = .834, r = -.33, 95% CI for proportion .50 = [.44, .60]) and was not correlated with the expected coordination success measure (ρ = -.060, p = .677).

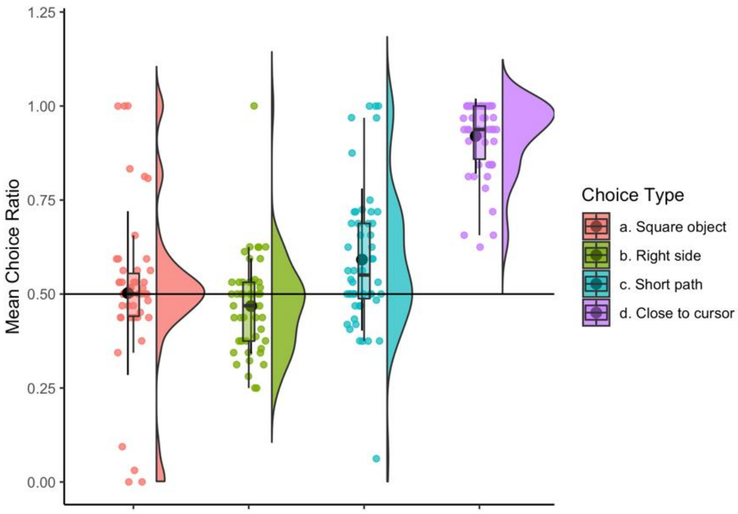

Figure 5.5: Proportions of (a) square shape choices, (b) right side object choices, (c) short path (relative to the participant's octagon) choices, (d) close to cursor choices in Experiment 1. In each box plot, black horizontal lines indicate medians, whiskers extend 1.5 times the interquartile range, black dots signal the means, with their whiskers extending to one SD; and the density plots illustrate the distributions of choice proportion data. Each colored dot represents a participant's choice ratio according to the given choice type. The horizontal line indicates chance level (.5).

'Side' bias. On average, regardless of its shape, participants chose the object positioned on the right side of the screen in 46.8% of the trials (SD = 0.13), which was a proportion not different from chance (Figure 5.5b, V = 263.5, p = .019, r = -.59, 95% CI for proportion .45 = [.40, .50]). We found a statistically significant relationship between expected coordination success and the side choices such that the lower the proportion of right-side choices, the higher the mean expected coordination success was (ϱ = -.508, p < .001).

This significant correlation may have been due to co-efficiency, because of the imbalance in the design where the co-efficient object appeared more often on the left (18 trials) than on the right side (14 trials): the proportion of left side choices (.53) was similar to the proportion of trials when the co-efficient object was on the left side (18/32 = .56). Therefore, the side of the screen where the square object was located was added as a categorical predictor to the regression models that predicted the probability of picking the square object (see Table B.7-12), in order to examine if the small, statistically not significant[46] side bias was due to co-efficiency or a coordination strategy in itself.

[46] N.B. the effect was statistically significant before Bonferroni-correction, therefore we thought it worth to add the Side factor to the regression models. The same applies to Experiment 2.

'Short path' bias. We also tested the ratio of participants choosing the object closer to their own starting octagon (i.e., the self-cost minimizing object) to see if there was any indication that participants were not trying to coordinate with their partner at all. On average, participants collected the black object positioned closer to their green octagon in 59% of the trials (SD = 0.19), a proportion higher than chance level (Figure 5.5c, V = 683, p = .001, r = .07, 95% CI for proportion .60 = [.52, .70]). This was a very small effect and unrelated to the mean expected coordination success (ϱ = .093, p = .521).

'Close to cursor' bias. We compared the ratio of choosing the object closer to the position of the participant's cursor at the beginning of the Decision phase to chance level. The average distance of the cursor from the chosen object was M = 83.6 pixels (SD = 100.7, see Figure B.1 in Appendix B). People overwhelmingly tended to choose the object located closer to their initial cursor position over the alternative object (Figure 5.5d, M = .92, SD = 0.10; V = 1275, p < .001, r = 1.00, 95% CI for proportion .94 = [.89, .97]). This strong tendency was not related to the mean coordination success of participants (Figure 5.4d, ϱ = -.083, p = .566).

5.3.2 Logistic Regression Models

Did participants simply aim to minimize their own individual action costs by picking the object nearest to their octagon instead of trying to coordinate? Alternatively, did they always choose the object that attracted their attention due to its proximity to their cursor, regardless of any of the cost disparities? Figure 5.6 shows the predicted choice probabilities according to the fitted cost-minimization models, before testing the data.

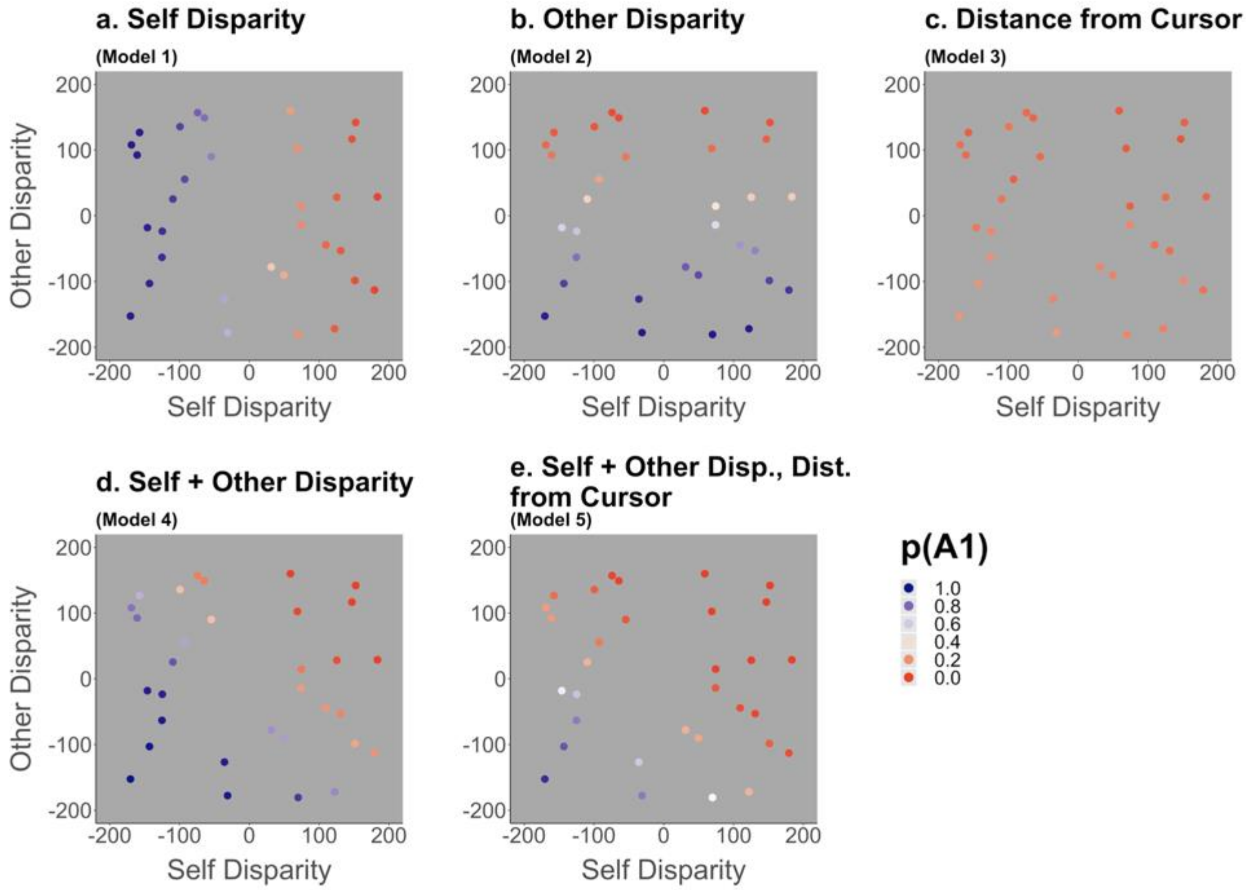

Figure 5.6: Predictions for optimal responses according to the (a) Self Disparity-, (b) Other Disparity-, (c) Distance from Cursor-, (d) Self and Other Disparity- and (e) Self and Other Disparity and Distance from Cursor-minimizing strategies, in the trials used in the study. The lower the disparity to be minimized according to a model, the higher the probability (blue) of picking the square/object A1. The predictions were calculated assuming a Boltzmann policy with the temperature parameter fixed to $k = 50$.

5.3.2.1 Primary models. We report the fixed-effect estimates for the five main mixed effects logistic regression models. Table 5.1 summarizes the maximum likelihood estimates (β coefficients) in log-odds for each model with 95% Confidence Intervals, with the predicted decreases in the odds of a square over a circle choice, due to a one pixel increase in the relevant predictor. Detailed tables of all models' estimates in odds ratios are included in Appendix B (Table B.1-15).

Models 1 and 2: Self Disparity, Other Disparity. To examine how much Self- and Other-cost minimization influenced decision-making, we fitted two models containing only the Self Disparity and the Other Disparity values as predictors. We found that the individual Self cost disparities of the participants had a statistically significant negative effect on the probability of them choosing the black square (Figure 5.7a, see Table 5.1 for the estimates, p = .002). We found a similar significant negative effect of Other Disparity on choices (Figure 5.7b, $p < .001$).

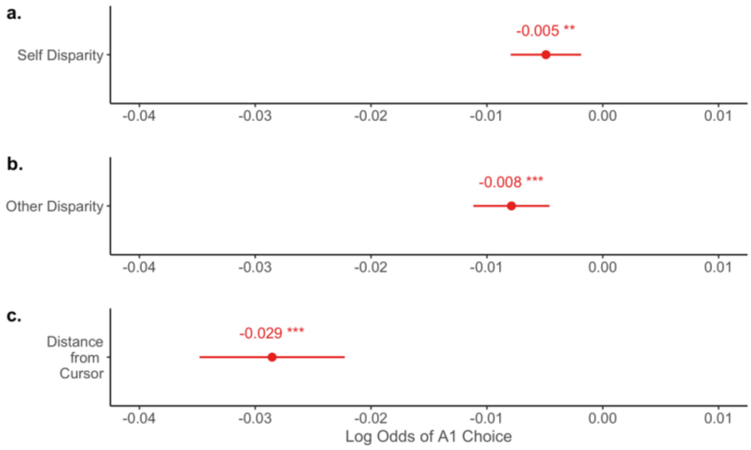

Figure 5.7: Log-odds of choosing the square object in Experiment 1, from the single-predictor mixed-effects logistic regression models on the effect of (a) Self Disparity, (b) Other Disparity and (c) Distance from Cursor. Subject ID was included as random grouping factor. Error bars show 95% confidence intervals (*** $p < .001$, ** $p < .05$).

Model 3: Distance from Cursor. The strong bias of participants to choose the object that was located closer to their cursor gave us reason to think that this factor by itself could have determined the participants' decisions. To test this hypothesis, we fitted a regression model containing only the Distance from Cursor as predictor. We found that this parameter had a statistically significant negative effect on the probability of participants choosing the square object (Figure 5.7c, Table 5.1, $p < .001$). This effect seemed to be larger than the effect of any of the two individual cost disparities' effect by themselves in Models 1 and 2.

Table 5.1: Parameter estimates from the primary fitted models in Experiment 1. Coefficients express the effect of the cost disparities in units of one pixel.

Object choice (A1) predictors	**β coefficient estimates Log-odds [95% CI], Standard Error**			**Odds Decrease (%) of A1 Choice due to a 1 px increase in predictor [95% CI]**
	β_{Self}	β_{Other}	$\beta_{DistanceFrom Cursor}$	
(1) Self Disparity	-0.005 [-0.008, -0.002], SE: 0.002			0.5 [0.2, 0.8]
(2) Other Disparity		-0.008 [-0.011, -0.005], SE: 0.002		0.8 [0.5, 1.1]
(3) Distance from Cursor			-0.029 [-0.035, -0.022], SE: 0.003	2.8 [2.2, 3.4]
(4) Self Disparity and Other Disparity	-0.008 [-0.012, -0.004], SE: 0.002	-0.011 [-0.015, -0.007], SE: 0.002		*Self:* 0.8 [0.4, 1.2], *Other:* 1.1 [0.7, 1.5]
(5) Self Disparity, Other Disparity, and Distance from Cursor	-0.010 [-0.014, -0.006], SE: 0.002	-0.008 [-0.012, -0.005], SE: 0.002	-0.031 [-0.039, -0.023], SE: 0.004	*Self:* 1 [0.6, 1.4], *Other:* 0.8 [0.5, 1.2], *Distance f. C.:* 3 [2.3, 3.9]

Models 4 and 5: Self and Other Disparity; Self and Other Disparity and Distance from Cursor. To directly assess the influence of co-efficiency on decision-making in the coordination task, we fitted multiple-predictor models using the linear combinations of Self and Other Disparity in Model 4 (expressing the effect of joint costs on choice), and the same combination with the added factor of Distance from Cursor in Model 5.

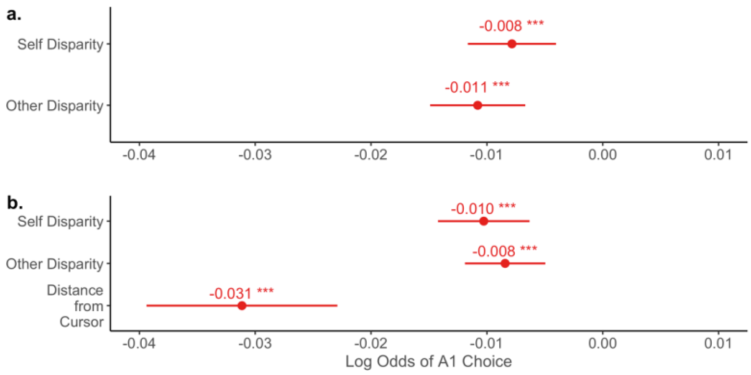

Figure 5.8: Log-odds of choosing the square object in Experiment 1, from the multiple-predictor logistic regression models on the effect of the linear combinations of (a) Self and Other Disparities, (b) Self and Other Disparities with Distance from Cursor. Subject ID was included as random grouping factor. Error bars show 95% confidence intervals (*** $p < .001$, ** $p < .05$).

In both models, every parameter exerted statistically significant negative effects (all *p*s < .001) on the probability of square object choices. In Model 4 (Figure 5.8a, Table 5.1), although the estimated effect of Other Disparity was numerically larger than that of the Self Disparity, the 95% CIs of the two estimates overlapped with one another. Therefore, we cannot conclude that there was a difference between how much the participants' object choices were influenced by the Self and Other Disparities.

Model 5 added the Distance from Cursor parameter to Model 4. The three estimates were statistically significantly different from zero in the expected negative direction (Figure 5.8b, Table 5.1, all *p*s < .001). The size of the two disparities' effects on the participants' square object choices were not different from one another (Figure B.2 shows the distribution of participants' β estimates, suggesting that both Self and Other Disparities were negatively weighted for almost everyone). The square object's initial Distance from Cursor exerted a substantially larger influence on the participants' decisions than the cost disparities.

5.3.2.2 Additional models. A simple categorization of the square as the co-efficient object[47] (Model 6) had a statistically significant effect on the probability of picking it (Log-odds = 1.25, 95% CI = [0.88, 1.61], $p < .001$). Participants were 3.48 times likelier (95% CI = [2.42, 5.02]) to choose the square when it was in the co-efficient position compared to when it was not. This odds ratio was not significantly larger than the odds ratio expected due to chance based on the congruent to incongruent trials' ratio in the experiment (75% / 25% = 3), which fell within the 95% CI.

The *Side* of the screen the square was positioned on did not have a significant effect on decisions by itself (Model 7: Log-odds = -0.29, 95% CI = [-0.62, 0.05], $p = .098$). The estimated log-odds were similar to what would be predicted by the odds of co-efficient objects appearing on each side of the screen (log((14/32) / (18/32)) = -0.25, see section 5.2.4 Design1). Further, *Side* in combination with other cost parameters (Models 8-12, Table B.8-12) did not have a systematic significant effect on decisions, and most of these models were not probable models of the data (see Table B.16).

5.3.2.3 Model Comparison. The R^2 values and Akaike weights for all estimated models (N = 15) are summarized in Table B.16. The proportion of variance explained by the fixed effects of the Self and Other Disparities was the largest in Model 5 ($R^2 = .741$), when Distance from Cursor was also added to the equation. The second best-fitting model in terms of variance explained by the factors of interest was Model 3, with Distance from Cursor only ($R^2 = .515$), whereas Self and Other Disparities by themselves (Model 4) explained a much smaller proportion of the variance ($R^2 = .154$). More than half of the variance was explained by the initial Distance from Cursor, but around another fifth was explained by the combined Self and Other Disparities, suggesting that the participants' behavior reflected a small effect of joint-cost minimization. The single-predictor Self and Other Disparity models fit the data particularly poorly ($R^2_{\text{Self}} = .047$, $R^2_{\text{Other}} = .098$).

[47] This is equivalent to basing decisions on the joint action costs, since a co-efficient object is by definition joint-cost minimizing, but the categorical variable does not differentiate between different magnitudes of the joint cost or the constituent individual action costs.

We found that Model 5 also seemed to be the relative best at predicting new data (AIC = 794.8, rounded w_5(AIC) = 1.00, Figure 5.9b shows the predicted square choice probabilities according to Model 5 against the observed data). Neither Model 4 (AIC = 1439.3, w_4(AIC) = 1.1^{e-140}), nor Model 3 (AIC = 1050.6, w_3(AIC) = 2.8^{e-56}) predicted data with similar accuracy (see Figure B.3 for plots of the observed data against their posterior predictions).

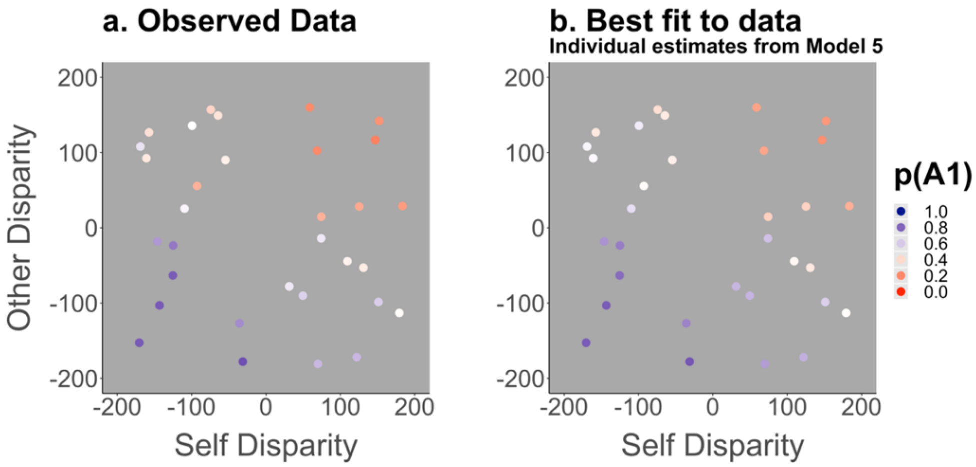

Figure 5.9: (a) Observed object A1 choice probabilities for all trials in Experiment 1 (N = 50). (b) Predicted choice probabilities based on the best-fitting model's mean of individual fixed effects estimates, including individual intercepts (Model 5, Self, Other Disparity and Distance from Cursor).

The second most probable model was Model 12 (Table B.12), which added the factor *Side* to Model 5's parameters This model had a probability of being a good predictor of future data at w_{12}(AIC) = 3.4^{e-05} (AIC = 815.4, marginal R^2 = .689).

5.3.3 Questionnaire data

We asked explorative questions about the identity of the partner to check if the belief manipulation of the remote experiment worked. We also asked the participants to report any conscious "strategy or decision rule" that they used in the task, to investigate if they were aware of making co-efficient choices, to the extent that they were doing so.

Beliefs about the identity of the co-actor. In Experiment 1, 14 participants reported that they believed their remote partner had been another human, 15 thought their partner was a computer, and 21 of them believed that they were the only player. We do not report estimations of regression models with this added Belief factor, since this was only a retrospective measure. However, the best-fitting model's individual fixed effects β estimates are plotted in Figure B.2 with colors based on reported belief, to provide some information of tendencies along this measure. The dispersion of the coefficients does not seem to reflect any particular belief-based differences in what cost disparity was minimized by the participants.

Self-reported strategy explanations. In general, the most often reported explanations[48] for the participants' behavior were picking the object (1) closer to the partner (n = 16), (2) closer to the participant (n = 11), (3) always picking the same shape (n = 11) and (4) picking the shape that was close to both the participant and the remote partner (n = 2). It seems that the majority of the participants thought they were making decisions based on individual movement costs: either (1) assuming that their partner would choose the object closer to themselves and the participants tried to match that expectation, or (2) not thinking about their partner and their potential cost-minimizing tendencies, and they simply aimed to pick the object always closer to themselves.

The participants with the highest proportions of co-efficient object choices (above 75%) gave the following answers: "I opted to click whichever shape that was closest to my partner, to try seeing things from their perspective"; "The object nearest to the octagon"; "Initially I thought there would be feedback so I was going to choose circles every time. Afterwards I chose whichever shape appeared closest to one of the two hexagons."; "*Tried to select closest objects overall for both players*"; "*I tried to pick what was closest to me and my partner.* I tried to see if there was some symmetry to my choices as well"; "Picked objects that were to the right of the octagon or closer"; "Nearest shape". These answers do not seem to suggest that the participants were aware that they were mostly choosing objects in the co-efficient position, though two of them gave answers suggestive of such a strategy (in italics).

[48] Other strategies cited were: a shape-based alternating pattern (n = 1), choosing the right-most shape (n = 1), choosing the object that was closest to one of the octagons (n = 1), choosing the object closer to the other half of the screen (n = 1), one answer was hard to interpret, and n = 9 participants did not answer. Some participants reported multiple rules, changing either conditionally or over time.

5.3.4 Discussion

In Experiment 1, we found that the probability of expected coordination success between the participants (56%) was higher than chance. This suggests that participants tried to match their object choices with their remote co-actors'. Their decision patterns, including an average co-efficient choice ratio of 63%, suggest a moderate effect of a strategy resembling the co-efficiency strategy. They also exhibited a small effect of picking objects close to their octagon. In 92% of the trials, the participants chose the object that was closest to their cursor's position at the beginning of a trial.

The decisions we observed were best predicted by the model incorporating a combination of similarly weighted Self and Other Disparities and the Distance from Cursor. Out of 50 participants, 46 people's beta coefficients were negative for both the Self and Other Disparities (Figure B.2), which indicates that the majority of the participants took into account both their own, and their partner's individual action costs when choosing objects. Without the cost disparities, the distance between the mouse and an object predicted decisions relatively poorly.

It is possible that a co-efficient configuration of objects acted as focal point solely on a visual level, without activating any representation of movement costs in the decision-maker. To test this hypothesis, we ran Experiment 2, where the participants were instructed to pick a black object by directly clicking on it only once. This eliminated the possibility of calculating joint movement costs, since the participant could not know where her partner's cursor was at the start of a trial.

5.4 Experiment 2 - Methods

5.4.1 Participants

The recruitment, testing procedure, and participant screening criteria on Prolific were identical to that of Experiment 1. Participants gave their informed consent and received monetary compensation through Prolific in exchange for their participation. The remuneration structure was also identical to Experiment 1. Fifty participants took part in Experiment 2 (23 males, age M =

31.8 years, *SD* = 10.7). Criteria for data inclusion were the same as for Experiment 1. We did not have to exclude any participants for not following the instructions.

5.4.2 Apparatus

The task was hosted on Gorilla and the participants used their own computers. Thirty-six participants used a computer mouse, 14 reported using a touchpad. The active screen area sizes during the task ranged from 776 x 388 pixels to 2543 x 1297 pixels.

5.4.3 Stimuli and Task

The same list of spatial layouts was used as in Experiment 1. The two experiments were identical except for the task instructions. Participants were now instructed to collect their chosen object by clicking on it only once, when their octagon turned green from orange. We emphasized that they should not click on any of the octagons. The position of the cursor at the beginning of each trial was not constrained by Gorilla. The timeouts and all other timing characteristics were also identical to those in Experiment 1 (see trial structure in Figure 5.3).

5.4.4 Design

The design was identical to Experiment 1's design (for details, see section 5.2.4 Design).

5.4.5 Procedure

The procedure was the same as in Experiment 1. The participants spent on average *M* = 14.7 minutes (*SD* = 3.18) on completing the task and the questionnaire on Gorilla.

5.4.6 Data Analysis

The same statistical analyses were conducted as in Experiment 1 (see section 5.2.6 Data Analysis, details of the estimated logistic regression models are available in Table B.17-32).

5.5 Results

5.5.1 Descriptive Statistics

5.5.1.1 Coordination Success. Overall, the ratio of matching decisions was statistically significantly higher than chance (Figure 5.10a, $M = 0.53$; $SD = 0.04$; $V = 1093$, $p < .001$, $r = .71$, 95% confidence interval (CI) for the proportion .53 = [.51, .55]).

5.5.1.2 Co-efficient choices. Overall, the participants chose the object in the co-efficient position 921 times out of 1590 trials[49] (Figure 5.10b, $M = .58$, $SD = 0.14$,). This was a proportion significantly lower than the chance level at .75 ($V = 50$, $p < .001$, $r = -.92$, 95% CI for the proportion .58 = [.52, .64]). This was a larger effect than in Experiment 1 ($r = -.88$), since the proportion was lower than before, but it was still significantly higher than 50% ($V = 796$, $p < .001$, $r = .25$, 95% CI for the proportion .59 = [.53, .66]). However, the proportion of co-efficient choices was even more strongly positively correlated with the expected coordination success of the participants than in Experiment 1 (Figure 5.10c, Spearman's $\varrho = .806$, $p < .001$), suggesting that the more often the participants chose the co-efficient object option, the likelier their choices were to match with others' in the sample.

[49] The proportions of choices were calculated based on trials with no missing choice values (i.e., where a valid decision was made: a square or circle was chosen on the participant's screen side). The data in Experiment 2 had 10 missing values in total, therefore the total number of trials was 1590.

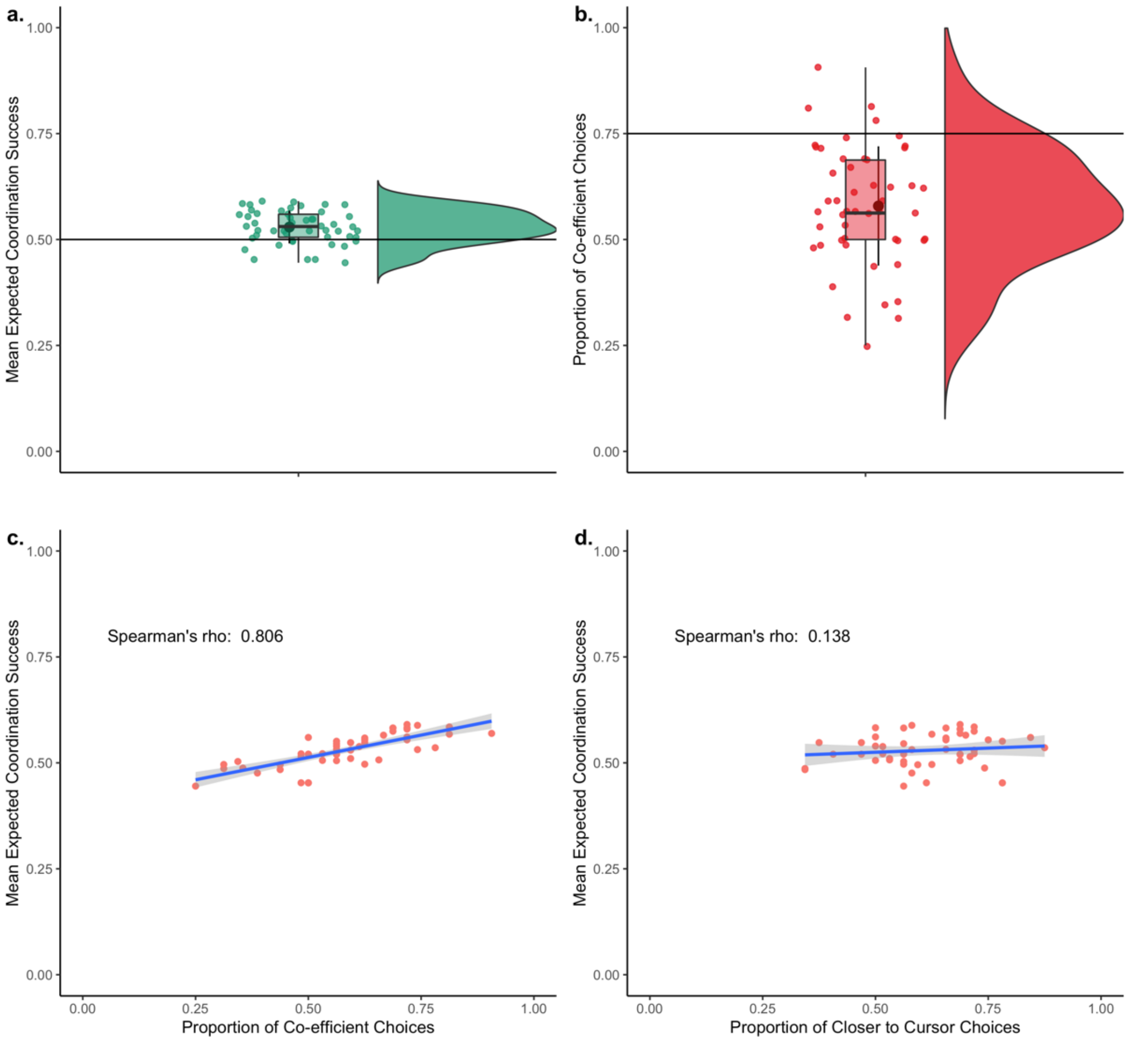

Figure 5.10: (a) Mean expected coordination success and (b) proportions of co-efficient object choices in Experiment 2. In each boxplot, black horizontal lines indicate medians, the bar's whiskers extend 1.5 times the interquartile range; black dots signal the means, with their whiskers extending to one SD; and the density plots illustrate the distributions of (a) mean coordination success and (b) choice proportion data. Each colored dot represents a participant's choice ratio for the given variable. The horizontal line indicates chance level (.5 for the mean expected coordination success and .75 for co-efficiency). Correlations between mean expected coordination success and (c) the ratio of co-efficient choices, and (d) the ratios of choosing the object closer to the mouse cursor at the beginning of a trial. The bold blue lines are fitted linear model predictions, the shaded area is a 95% confidence interval.

5.5.1.3 Object Shape, Side, Short Path Biases. *'Object shape' bias.* The participants chose the square object on average 17 times out of 32 trials (Figure 5.11a, M = .53, SD = 0.23), which was a proportion not different from chance (V = 496.5, p = .577, r = -.22, 95% CI for proportion .52 = [.45, .64]). This choice proportion was not correlated with the expected coordination success measure (ϱ = .143, p = .321).

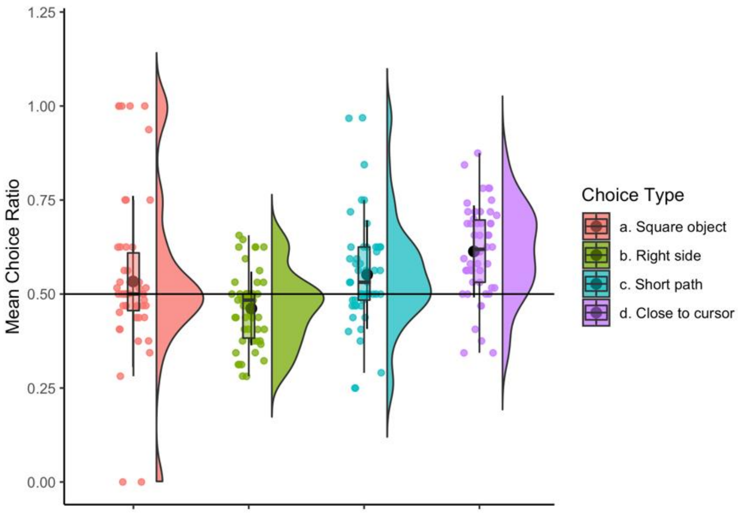

Figure 5.11: Proportions of (a) square shape choices, (b) right side object choices, (c) short path (relative to the participant's octagon) choices, (d) close to cursor choices in Experiment 2. In each data bar, black horizontal lines indicate medians, whiskers extend 1.5 times the interquartile range, black dots signal the means, with their whiskers extending to one SD; and the density plots illustrate the distributions of choice proportion data. Each colored dot represents a participant's choice ratio for the given choice type. The horizontal line indicates chance level (.5).

'Side' bias. The participants chose the object positioned on the right side of their screen on average 14.7 times out of 32 trials, which was a proportion only numerically lower[50] than chance (Figure 5.11b, $M = .46$, $SD = 0.10$, $V = 221$, $p = .011$, $r = -.65$, 95% CI for proportion .45 = [.40, .50]). The correlation between the proportion of right-side choices and expected coordination success was not statistically significant ($\alpha = .01$; $\varrho = -.290$, $p = .041$). Side was again added to the additional regression models to explore if the small bias was due to co-efficiency or an independent coordination strategy.

'Short path' bias. The participants chose the object relatively closer to their octagon on average 17.5 times, a proportion marginally significantly different from chance (Figure 5.11c, $M = .55$, $SD = 0.14$, $V = 656.5$, $p = .010$, $r = .03$, 95% CI for the proportion .55 = [.50, .62]). This moderately correlated with the expected coordination success ($\varrho = .380$, $p = .007$).

'Close to cursor' bias. The average distance of the cursor from the chosen object at the start of a trial was $M = 162.4$ pixels ($SD = 137.5$; see Figure B.4). The proportion of object choices that were relatively closer to the cursor at the start of the Decision phase was statistically significantly

[50] As in Experiment 1, this effect was statistically significant before Bonferroni-correction.

higher than chance. This effect was medium-sized, not as strong as in Experiment 1 (Figure 5.11d, M = .61, SD = 0.12, V = 986.5, p < .001, r = .55, 95% CI for the proportion .63 = [.57, .68]). The tendency to collect the object closer to the cursor was not related to the level of expected coordination success (Figure 5.10d, ϱ = .138, p = .341).

5.5.2 Logistic Regression Models

5.5.2.1 Primary models. *Models 1 and 2: Self Disparity, Other Disparity.* According to the single-predictor models (Figure 5.12a-b, Table B.17-18) both Self and Other Disparity by themselves had statistically significant (Self: p = .012; Other: p < .001), negative effects on the probability of square object choices. The estimates are summarized in Table 5.2, together with their predicted % decreases in the odds of a square choice over a circle choice.

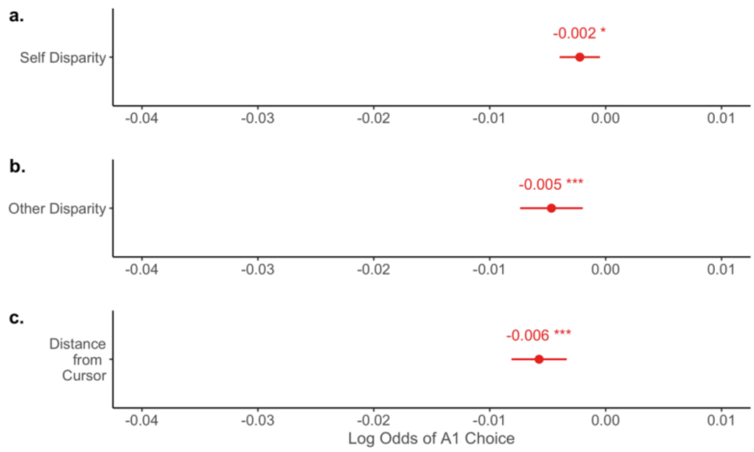

Figure 5.12: Log-odds of choosing the square object in Experiment 2, from the single-predictor mixed-effects logistic regression models on the effect of (a) Self Disparity, (b) Other Disparity and (c) Distance from Cursor. Subject ID was included as random grouping factor. Error bars show 95% confidence intervals (*** p < .001, ** p < .05).

Table 5.2: Parameter estimates from the primary fitted models in Experiment 2. Coefficients express the effect of the cost disparities in units of one pixel.

Object choice (A1) predictors	β coefficient estimates Log-odds [95% CI], Standard Error			Odds Decrease (%) of A1 Choice due to a 1 px increase in predictor [95% CI]
	β_{Self}	β_{Other}	$\beta_{DistanceFromCursor}$	
(1) Self Disparity	-0.002 [-0.004, -0.0004], SE: 0.001			0.2 [0.05, 0.4]
(2) Other Disparity		-0.005 [-0.007, -0.002], SE: 0.001		0.5 [0.2, 0.7]
(3) Distance from Cursor			-0.006 [-0.008, -0.003], SE: 0.001	0.6 [0.3, 0.8]
(4) Self Disparity and Other Disparity	-0.004 [-0.006, -0.001], SE: 0.001	-0.006 [-0.009, -0.003], SE: 0.002		Self: 0.4 [0.1, 0.6], Other: 0.6 [0.3, 0.9]
(5) Self Disparity, Other Disparity, and Distance from Cursor	-0.003 [-0.005, -0.001], SE: 0.001	-0.006 [-0.010, -0.003], SE: 0.002	-0.005 [-0.008, -0.003], SE: 0.001	Self: 0.3 [0.1, 0.5], Other: 0.6 [0.3, 1.0], Distance f. C.: 0.5 [0.3, 0.8]

Model 3: Distance from Cursor. The effect of an object's distance from the participant's cursor's position at the start of a trial on choices was statistically significant and negative, as expected (Figure 5.12c, Table 5.2, $p < .001$). In the three single-predictor models, the sizes of the estimated effects

of the parameters were not likely to differ from one another, as their 95% confidence intervals overlapped.

Models 4 and 5: Self and Other Disparity; Self and Other Disparity and Distance from Cursor. We analyzed the effect of the linear combinations of Self and Other Disparities in Model 4. Self and Other cost disparities by themselves both had statistically significant negative effects on the probability of A1 choices, of similar magnitudes (Figure 5.13a, Self Disparity: p = .001; Other Disparity: p < .001).

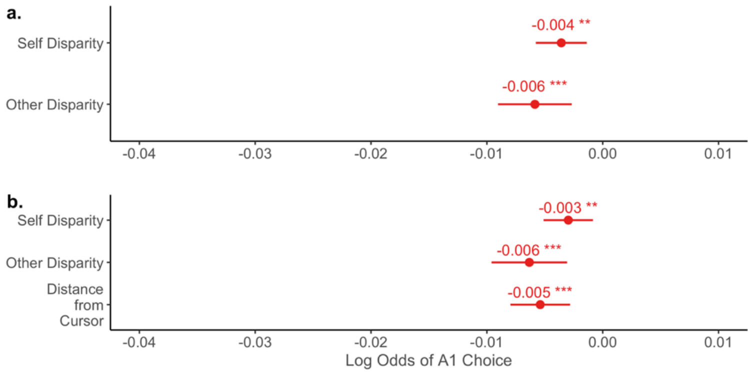

Figure 5.13: Log-odds of choosing the square object in Experiment 2, from the multiple-predictor logistic regression models on the effect of the linear combinations of (a) Self and Other Disparities, (b) Self and Other Disparities with Distance from Cursor. Subject ID was included as random grouping factor. Error bars show 95% confidence intervals (*** p < .001, ** p < .05).

The addition of Distance to Cursor to Model 4 resulted in similar effects in Model 5. All three of the predictors had statistically significant (Self Disparity: p = .006; Other Disparity and Distance from Cursor: both ps < .001) negative effects on the square choice probability, and the sizes of the effects were comparable to one another (Figure 5.13b, Table 5.2). In summary, both Model 4 and 5's estimates reflect a strategy resembling joint-cost minimization, as the weights on the Self and Other Disparities were similar in both models. We also found similar-sized effects on decision-making by the cost disparity parameters and by the distance between a participant's cursor at the start of the Decision phase. This is different from Experiment 1 where the size of the latter factor's effect outweighed the two disparities' effects.

5.5.2.2 Additional models. The categorical ID of the square as the co-efficient object in a trial (Model 6, Table B.22) had a statistically significant effect on the choice probability, though the effect was not as large as in Experiment 1 (Log-odds = 0.79, 95% CI = [0.40, 1.17], p < .001). Participants were 2.2 times likelier (95% CI = [1.5, 3.23]) to choose the square when it was in the co-efficient position, rather than in a jointly sub-efficient position. This odds ratio was not different from the odds ratio expected due to chance based on trial congruency (OR = 3).

The *Side* of the screen that the square was positioned on was found to have influenced object choices by itself (Model 7: Log-odds = -0.35, 95% CI = [-0.61, -0.09], p = .009). However, the log-odds of the co-efficient object appearing on the right side of the screen (-0.25) was within the estimated 95% CI for the log-odds of a right-side choice. Models using the combinations of Side and different cost parameters provide further indication whether cost disparities or Side influenced the probability of square choices to a greater extent.

The results of Model 11 (Table B.27), which included the combination of Self and Other cost disparities together with Side, found that only the two cost disparities had significant main effects on choice probability (Self Disparity: Log-odds = -0.004, 95% CI = [-0.007, -0.002], p = .001; Other Disparity: Log-odds = -0.005, 95% CI = [-0.008, -0.001], p = .006). Neither any of the interaction effects, nor the main effect of Side on square choices were statistically significant (Side: Log-odds = -0.286, 95% CI = [-0.596, 0.024], p = .070). Model 12 (Distance from Cursor was added to Model 11, Table B.28) further confirmed the lack of Side effect, independent of the effects of cost disparities and cursor distance[51].

5.5.2.5 Model Comparison. Table B.32 summarizes the R^2 values and Akaike weights for every model (N = 15). The proportion of variance explained by the fixed effects was the largest in Model 12 (marginal R^2 = .166), a version of Model 5 extended with the *Side* factor. The second best-fitting model was Model 5, with the predictors Self and Other Disparity with Distance from Cursor (R^2 = .148; Figure 5.14b shows the predicted square choice probabilities based on Model 5's estimates against the observed data). Distance from Cursor by itself explained a much smaller

[51] N.B. the correlation between the sum of the Self and Other disparities (Joint Disparity) and the Side factor was small and not statistically significant (ϱ = .048, p = .054).

proportion of variance (Model 3, R^2 = .087), just like Self and Other Disparities by themselves (Model 4, R^2 = .059). As in Experiment 1, the single-predictor Self and Other Disparity models fit the data particularly poorly (R^2_{Self} = .012; R^2_{Other} = .037). The observed behavior was best described by models incorporating Self and Other Disparities together with Distance from Cursor.

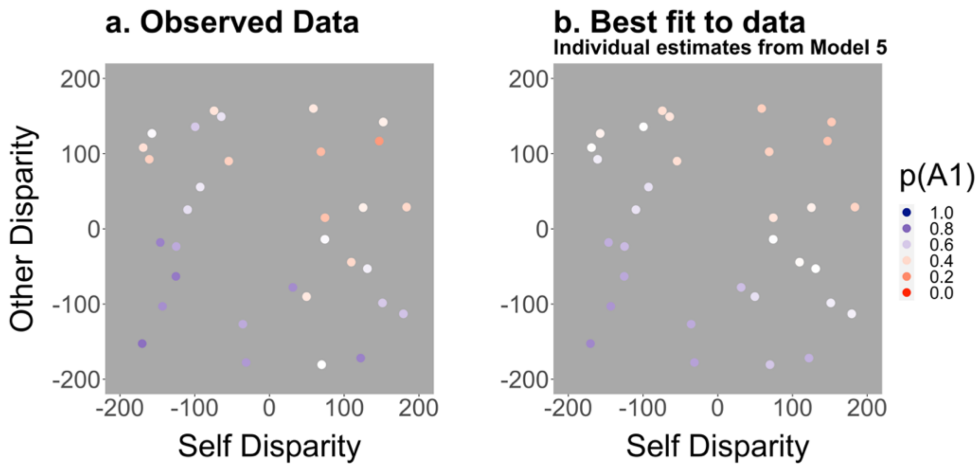

Figure 5.14: (a) Observed object A1 choice probabilities for all trials in Experiment 2 (N = 50). (b) Predicted choice probabilities based on the best-fitting model's mean of individual fixed effects estimates, including individual intercepts.

Regarding predictive accuracy, we found that Model 5 was again expected to be the relative best at predicting new data (AIC =1685.5, w_5(AIC) = .96). Model 4 using Self and Other Disparity only (AIC =1729.1, w_4(AIC) = 3.2^{e-10}), and Model 3 with Distance from Cursor only (AIC = 1890.9, w_3(AIC) = 2.4^{e-45}) were both estimated to predict data with much lower accuracy than the best-fitting model (Figure B.6 shows the observed data against the posterior predictions of Models 3 and 4).

5.5.3 Questionnaire data

Beliefs about the identity of the co-actor. Eighteen participants reported that they thought their remote partner was another human, 15 thought their partner was a computer, and 15 of them believed that they played on their own (two participants did not answer). Figure B.5 shows the estimated beta coefficients for the Self and Other Disparities in the best model of the behavior (Model 5), with added colors for each participant's reported belief about the partner. A visual inspection suggests that the weights on the action cost parameters were not clustered based on belief.

Self-reported strategy explanations. Most often, participants reported that they chose the object that was closest to their partner (*n* = 12, six of whom specifically mentioned the partner's octagon), or to themselves (*n* = 10, four of whom mentioned their octagon). Twelve participants reported that they always picked the same shape, while some other participants followed a shape-based alternating pattern either in part of, or the whole task (*n* = 3). Three participants chose objects always on the same side (right: *n* = 2, left: *n* = 1). One participant reported that they picked the object that was closer both to them and their partner, whereas another participant picked the objects closer to the octagon on either side of the screen[52]. One participant made specific reference to the fixation cross and their line of sight at the beginning of a trial, stating that they collected the object that was closer to where their gaze was fixated before a trial's start, and also to where an object had been picked in the previous trial.

Were those participants who made the highest proportions of co-efficient object choices consciously choosing joint-cost minimizing objects? The answers by participants with more than 75% co-efficient choices were: *"I selected the closest shape to me, except when my partner's shape was exceedingly close to them."; "I was choosing the shape that was closest to the octagon on either side";* "Closest to the Octagon";" That my partner might choose the shape closest to their hexagon so I matched my choices accordingly". Half of these answers suggest that the participants were consciously making choices that took into account both players 'movement distances (in italics). More importantly, all of these make reference to the octagons that in this experiment were *not* part of the object collection procedure. In total, 11 answers mentioned the octagon(s). This suggests that beyond the octagon's

[52] One person reported that they chose the object closer to the middle of the screen, and two answers were hard to interpret ("chose the closest symbol in their box"; "try to select the mirror image"). Fourteen participants did not answer the strategy question. Some participants gave compound answers, describing how they changed their strategies over time.

role as a trigger for starting to move, at least fifth of the participants viewed the octagon(s) as relevant to their decision-making.

5.5.4 Comparison of the two experiments

We compared the main findings of the two experiments directly. The expected coordination success was statistically significantly larger in Experiment 1 (M = .56, SD = .06) than in Experiment 2, a medium-sized effect (M = .53, SD = .04; two-tailed[53] Mann-Whitney test: U = 1703.5, p = .002, 95% CI for the difference in median proportions .04 = [.02, .06], r = -.36). We found a marginally significant difference between the co-efficient choice proportions in the two experiments (Experiment 1: M = .63, SD = .14; Experiment 2: M = .58, SD = .14; Mann-Whitney test: U = 1532.5, p = .051, 95% CI for the difference in median proportions .06 = [-5.5^{e-06}, 1.2e-01], r = -.23).

The best-fitting logistic regression model using Self and Other Disparity with Distance from Cursor was estimated again on data collapsed over experiments (see Table B.33-35), both with and without the added binary predictor *Experiment*. We found that the model including *Experiment* fit the data better and was expected to be better at out-of-sample prediction than the model without it (with: AIC = 2514.8, w_{5_Exp}(AIC) = 1.00, marginal R^2 = .349; without: AIC = 2568.4, w_5(AIC) = 2.3e-12, R^2 = .259). The Experiment factor interacted with Self Disparity (X Experiment interaction Log-odds: 0.006, 95% CI = [0.002, 0.10], p = .004; β_{Self_Exp1} = -0.009, 95% CI = [-0.012, -0.006]; β_{Self_Exp2} = -0.003, 95% CI = [-0.006, -0.0003]) and Distance from Cursor (X Experiment interaction Log-odds: 0.020, 95% CI = [0.013, 0.026], p < .001; β_{Dist_Exp1} = -0.025, 95% CI = [-0.030, -0.020]; β_{Dist_Exp2} = -0.006, 95% CI = [-0.010, -0.001]). These predictors' effects on the probability of square choices were stronger in Experiment 1 than in Experiment 2 (Figure B.7). The effect of Other Disparity on choices did not differ across experiments.

[53] We report the results of two-tailed Mann-Whitney tests for better interpretability since one-tailed tests yielded 95% CIs with infinite bounds. The conclusions based on p-values obtained are the same regardless of test type.

5.5.5 Discussion

We found in Experiment 2 that the level of mean expected coordination success was higher than chance (although lower than in Experiment 1), and it strongly correlated with the proportion of object choices that were in a co-efficient position. We also replicated the effect of proximity of the cursor to an object at the start of a trial: though less strongly than in Experiment 1, the participants tended to choose the objects that were closer to their cursor. This is essentially self-cost minimization under the current "one-click" instructions (even though octagon-based "Self costs" seemed also to be minimized).

Importantly, regression analyses suggested that the decisions observed in Experiment 2 were best described by the same model as in Experiment 1: the linear combination of similarly weighed Self and Other cost Disparities and an object's Distance from Cursor. The effects of Self Disparity and Distance from Cursor were smaller than in the first experiment. Based on the answers provided in the post-task questionnaire, it appears that the participants assigned relevance in their decision-making to the two octagons beyond what we intended. They often referred to the octagons as anchors for the actions – both their own and their partners' –, even though the octagons were not included to be clicked on in this version of the task.

5.6 General Discussion

Coordinating in uncertain situations has long been in the focus of investigation by classical game theory and alternative approaches such as level-k or team reasoning theories. The latter theory postulates that strategy pairs that maximize the joint payoffs of a group can serve as focal points, which team-reasoning agents are able to identify and select. This often ensures successful coordination between people (Bacharach, 2006). Here, we took steps to extend the investigation of focal points in coordination to movement-related decision-making, by testing if people use co-efficiency in joint actions as a payoff-related focal point when they do not have information about their partners' decisions in repeated coordination games.

We designed an online version of the object matching task from Study 2 (Chapter 3), where co-actors were instructed to collect matching black objects with one another on a shared screen. In each trial of the task in Experiment 1, participants chose between a black square or a black circle and collected it by clicking first on their designated starting octagon, then on the chosen object, and finally, on their octagon again. This three-step procedure was identical to the one in our lab-based study in Chapter 3. In Experiment 2, we changed the instruction and asked participants to collect objects by simply clicking on them once. In both experiments, the starting position of the participant's mouse cursor was not fixed but under the participant's control. The participants were told that their remote partner was playing along with them, trying to pick a matching black object on their side of the screen. Crucially, the participants did not receive feedback about their partner's decisions and were told that they would score points for each object choice that matched with their partner's.

We analyzed object choices regarding their co-efficiency (i.e., whether the chosen object was part of the joint-cost minimizing pair of objects), and other potentially salient features such as shape, relative distance from the participant's own octagon, the side of the screen on which it was positioned, and its distance from the cursor at the start of a trial in pixels (i.e., if the chosen object was the object relatively closer or farther from the cursor). In both experiments, we found that the proportion of co-efficient object choices was significantly higher than 50% – a moderate effect in Experiment 1, and a small effect in Experiment 2. In Experiment 1, we found a bimodal distribution of both co-efficient decisions and expected coordination success, showing that some people tended to choose the co-efficient object, whilst others did not, and some people had a higher than chance probability of coordination success with others, whilst others were at chance. We did not observe such bimodality in Experiment 2.

These results indicate that some participants tended to choose objects that were part of the co-efficient solution over the sub-efficient alternative. However, the analyses of these proportions in both experiments suggested that co-efficient choices could have also been a result of a Self or

Other-cost minimizing strategy, since the average choice proportions were lower than 75% (the incidence of congruent trials in the experiments). To rule out such alternative strategies, we predicted decisions in logistic regression models by different combinations of cost disparities. Nevertheless, the overall co-efficient choice proportions were strongly correlated with a mean expected coordination success measure in both experiments, which suggested that the more often the participants chose objects in co-efficient pairs for whatever reason, the likelier they were to match with the decisions of the rest of the sample.

Participants could have chosen arbitrary characteristics like shape or left-right position in a layout, and by consistently selecting for that characteristic, they could have expected an overall coordination success of around 50% with their partner (provided they assumed that their partner also chose and persevered on a selected feature of the objects). However, we did not find convincing evidence for alternative strategies that could help improve coordination success: in neither of the experiments did participants show a significant shape or screen side bias.

Notably, we found a large (Experiment 1) and a medium-sized (Experiment 2) effect of the cursor's initial distance from an object. This factor strongly influenced decisions, however, by itself it did not fully explain behavior. Regression analyses confirmed our main hypothesis by revealing that in both experiments, the models that included similarly weighed Self and Other cost disparities in addition to Distance from Cursor predicted observed data the best. As Chapter 3 suggested, an equally weighed linear combination of Self and Other disparities formalizes a joint-cost minimizing, co-efficient strategy when people interact in real time. Based on these results, we conclude that objects that were part of object pairs that would minimize joint action costs in a real time coordination task seemed to help coordination to some extent in the present study, as focal points.

We ran Experiment 2 to test the hypothesis that the moderate effect of co-efficiency in Experiment 1 was related to movement costs, and not to the visual arrangement of the object layout. We hypothesized that if Experiment 2 yielded results similar to Experiment 1, that could be taken as evidence that the co-efficiency effect we observed was likely visual, not motor based.

A priori, the participants were expected not to take into account the octagons in their planning (beyond their role as a trigger for starting object collection at the start of a trial) and not to calculate joint action costs. The results, however, paint a different picture.

Despite the one-step object collections procedure, at least 20% of the participants in Experiment 2 mentioned the octagons in their strategy reports in the post-task questionnaire.[54] Furthermore, the effect on decision-making of the partner's costs as defined by Other disparity reflects that most participants assumed that their remote partner always started their movement from their own octagon, despite knowing that they both followed the same task instructions. The qualitative similarity of the main results across the two experiments suggests that the co-efficient object pairs tended to be recognized as focal points for coordination by some people even when they were not required to perform (their half of) the action sequences that would have been optimized by co-efficient choices. That is, even in the case of "imagined" joint actions where the joint costs were not actually invested, people sometimes used co-efficiency as focal points, although this effect was smaller than in the case of the "real" action sequences in the first experiment. This raises the possibility that co-efficiency in our task could act as an abstract rule to select focal points for coordination in each trial.

We do not wish to make claims about how our participants identified focal points, whether team reasoning (Bacharach, 2006; Sugden, 1993) or other mechanisms like reasoning in a cognitive hierarchy (Camerer et al., 2004) helped participants to "match" their decisions with their remote partner's decisions above chance level. Rather than arguing for the primacy of one mode of reasoning over the other, the literature suggests that people engage in both, depending on the features of a coordination game (Bardsley et al., 2010; Faillo et al., 2017). Participants in our study could have searched for a selection rule that best distinguishes black objects from each other, thus arriving at the co-efficient configurations of objects; or they could have reasoned that their partner

[54] Omitting the octagons was not an option as that would have changed the relative salience of visual features between the two experiments, introducing a confound. In the instructions for Experiment 2, we emphasized that choosing the octagon was not a valid choice for coordination, only the black objects.

might reason on a dyadic level and try to minimize the total distances to objects, and therefore participants could have best-responded to this possibility. Although the former possibility seems more parsimonious, moving beyond the realm of speculation would necessitate further experiments that address the reasoning behind coordination. Similarly, questions related to the effect of the (presumed) agency of a partner would best be assessed by experiments specifically tailored to them.

To summarize, our study provided some evidence for the benefit of using co-efficiency as a decision rule in coordination games that involve real or imagined movements. Selection rules were proposed by Schelling (1960) and consequently tested by Mehta et al. (1994a) as a type of salience ("Schelling salience"). Co-efficient object pairs seem to be sometimes recognized as focal points on which players can converge and match their remote partner better than if they could match by using a random selection strategy. It might be by focusing on "what would we do, if we acted together now?" (cf. Bacharach, 2006), that participants coordinated their choices more often than 50%. The type of reasoning behind the effect we found remains to be explored.

Chapter 6. General Discussion

The aim of this dissertation was to empirically investigate whether and how the normative principle of rationality, that guides the planning and production of individual actions and enables the understanding of others' observed actions, applies to social contexts. Specifically, I investigated whether a principle of rational joint action might realistically describe behavior in joint actions, and if it can help coordination under uncertainty. The rationality principle postulates that agents' actions serve to bring about desired goals, and they do so by the most efficient available means under given environmental constraints (Gergely & Csibra, 2003). Previous research supports the idea that the rationality principle has its basis not only in animal, but also in human behavior, since in individual action contexts, people often act in optimal ways given the costs and rewards associated with potential actions and their outcomes (Todorov, 2004; Trommershäuser et al., 2003a, 2003b; Wolpert & Landy, 2012). Furthermore, according to some studies, cooperative actions may also be understood based on the principle of rational action (Gredebäck & Melinder, 2014; Mascaro & Csibra, 2014), since infants expect that these actions unfold jointly efficiently – even if this might sometimes require actors to act individually inefficiently. I argued that sharing a joint goal in cooperative activities justifies the expectation that co-actors should also share the costs and benefits related to the attainment of the goal, and therefore co-actors ought to choose joint action plans that achieve these goals maximally efficiently. In essence, if we take individual actions to maximize the expected utility of individuals, then by extension, joint actions should maximize the expected utility of a collective agent (Gilbert, 2006). Is there a behavioral basis for such an expectation about social interactions?

In this dissertation, I hypothesized that people plan joint actions that conform to the principle of rational joint action, that is, actors choose the most co-efficient (jointly efficient) means of achieving a shared goal. In the experiments I presented, we operationalized action costs as proportional to distance. Following an initial study that confirmed this hypothesis, we tested a hypothesis concerning the computation of joint costs. Then we addressed the question whether co-efficient decision-making can be generalized to situations where people perform joint action sequences composed of different kinds of actions. Finally, we tested whether an expectation of co-efficiency helps in coordination problems by providing a focal point for actors who have to guess their remote interaction partner's decisions. In the following, I summarize the findings of these four empirical studies and discuss their theoretical implications, as well as the questions they raise for future research.

6.1 Maximizing co-efficiency in coordination

In Chapter 2, we investigated how people distribute the costs of a joint action sequence between themselves and a co-actor. We hypothesized that a decision-making actor would maximize the co-efficiency of the dyad by choosing an action plan that minimizes the overall costs of an action sequence, given the available options. Previously, several joint action studies had found that people tend to modulate their own actions to reduce the effort the partner would have to expend to reach the dyad's shared goal (e.g., Dötsch & Schubö, 2015; Meyer et al., 2013; Ray & Welsh, 2011). A potential explanation for such facilitatory behaviors is the shared-effort model proposed by Santamaria and Rosenbaum (2011). This model postulated that people coordinate their actions with others if the shared effort of the coordination would be lower than the sum of the costs that the individuals would otherwise incur, were they to act on their own. Sharing beliefs about such implicit cost comparisons was assumed to be a precondition (Santamaria & Rosenbaum, 2011).

Like the aforementioned joint actions studies, Santamaria and Rosenbaum's (2011) observational study found that people willingly incurred effort to help another person by holding a door open for them. However, this in itself did not provide proof for minimizing shared effort. It is possible that by investing effort themselves, people aimed to reduce the effort (action costs) of the other person only, rather than taking into account the sum of all the actors' costs in the given situation.

Chapter 2 provided robust evidence that people took into account the total efficiency of joint action sequences, and did not aim to minimize only their own, or their partner's individual action costs. In an experiment on individual action planning, we confirmed efficiency maximization (in line with previous results on optimal action planning, see section 1.2 Action planning as rational decision-making), which was paired with a tendency to complete a sub-goal earlier, rather than later, when the efficiency of the whole action was unaffected by decisions. This is a result similar to findings on precrastination, as described by Rosenbaum and colleagues (2014) and subsequent research (e.g., Fournier, Coder et al., 2019). Although we did not probe this individual decision-making pattern further, it serves as a reminder that cognitive effort is part and parcel of most cost-benefit analyses when people decide about how to act (Kool et al., 2010), and they are traded off against movement effort (Feghhi & Rosenbaum, 2019).

In three joint experiments, the same two-part actions as in the individual experiment were distributed across co-actors. We consistently found an effect of the total path length on people's choices: when co-efficiency could be minimized (i.e., the alternative total path lengths were unequal), the participants adapted their decisions to the particular action plan options. When taking over the partner's effort benefitted co-efficiency, people were inclined to do so. However, when *not* taking over the partner's effort was co-efficient, they tended to choose the path that was individually more efficient for themselves. The co-efficiency effect was slightly larger when it coincided with reducing the co-actor's path length; when co-efficiency could not be minimized, the participants were biased to reduce the co-actor's individual costs. These findings support the interpretation that in the joint action context, the actors took into account both their own and their partners' individual action costs, and likely aimed to minimize their sum.

On the one hand, part of our results conceptually replicated previous findings on facilitation in joint actions. The studies that found facilitation interpreted it as proof for task co-representation and planning with a partner in mind. The behavior that we observed in neutral trials, where we found a reliable tendency for people to make decisions that shortened the distance assigned to a co-actor, are most directly comparable to findings by Scharoun and colleagues (2017) and Gonzalez and colleagues (2011). They found that people placed objects closer to their partner, and in orientations that suggested a consideration for the partner's anticipated effort in the second part of the joint sequence. Gonzalez et al. (2011) interpreted this as an extension of the first actor's end-state comfort to the second actor's beginning-state comfort, placing similar importance on the two individuals' respective efforts. Other kinds of facilitatory behavior like additional rotations scaled to a co-actor's task (Constable et al., 2016; Dötsch & Schubö, 2015) are also consistent with our findings, demonstrating that a variety of action modulations may be used to realize cost reduction for a partner.

On the other hand, our findings extend this joint action literature on planning and the co-representation of task constraints, by also showing *non-facilitatory* behavior when that benefitted the dyad. We provided confirmation for the authors' speculation that the shared-effort model (Santamaria & Rosenbaum, 2011) may be behind the facilitatory behaviors observed. We contrasted this hypothesis with the alternative that people simply try to altruistically reduce a co-actor's costs, through the use of experimental trials where altruistic facilitation would hinder shared effort reduction (co-efficiency).

This study raised a host of questions for future research, some of which we subsequently addressed in this dissertation (see sections 6.2 Computing joint action costs and 6.3 Integrating the costs of different types of actions), but many still remain to be explored. For example, starting from this study and throughout this entire project, we manipulated utilities by systematically manipulating only action costs while fixing the rewards at the completion of a trial. If people calculate and compare expected utilities for the dyad (i.e., joint reward minus joint cost) to make jointly rational and co-efficient decisions, then a manipulation of the joint reward should change the probabilities of making co-efficient decisions in predictable ways (see Le Bars et al.'s [2020] preprint for an example of manipulating both rewards and costly motor noise in a coordination task). Alternatively, it is possible that people are unable or unwilling to estimate full expected utilities and choose between action plans based only on either part of the utility. This can be easily tested by introducing different rewards for each trial in a task similar to the one in Chapter 2 and would provide further insight into joint action planning processes.

Additionally, the boundary conditions of co-efficiency promise to provide an exciting further direction for research. For instance, although Chapter 2 suggests that the reciprocity of decision-making is not a condition for co-efficient decisions, it is possible that the previous (un)cooperativeness of a co-actor could influence the propensity to prioritize joint-cost minimization. Scharoun and colleagues' (2017) helpfulness manipulation design could provide an example for a simple method to test the hypothesis that people might regard co-efficient decisions as an expression of cooperativeness. After a history of not co-efficient (e.g., consistently selfish) behavior from a co-actor, they might reduce cooperative behavior by making decisions that help reach the joint goal, but not in the most co-efficient way.

6.2 Computing joint action costs

Chapter 3 focused on the question of how people estimate the joint costs of an action sequence. We tested the hypothesis that when the costs of individual co-actors' actions are on the same scale (distance), people compute the potential joint action costs as a weighted sum of the co-actors' individual action costs. Although in Chapter 2, the different findings between individual and joint contexts of the neutral trials suggested that in the joint task, participants took into account the actors' individual path lengths to minimize joint costs, we could not confidently conclude this. This is because the movement sequences that a decision-making actor was presented with were quasi-continuous (from one corner of the screen to the other), and the actor could have planned the action as if they would be performing it alone; and could have chosen a path based on such an individualistic plan. To test whether the costs of the individual actions were summed to estimate the joint costs, in Chapter 3, we used an object matching task where each action sequence was composed of non-connecting movements. The separation of individual actions enabled the independent parametric manipulation of individual and joint action costs. Using multiple regression models, we could tease apart the contributions of each type of cost to the decision-making.

Our findings from three experiments showed that the participants chose objects and corresponding action plans that minimized the combination of expected Self and Other action costs. They made such decisions after a brief initial phase (a few trials) of mainly minimizing their own individual action costs. We found that when choosing between objects to collect, the participants weighted their own costs roughly equally to their partners' costs: the population-level relative weights were .56 (95% HDI = [.42, .71]) for the Self, and .44 (95% HDI = [.29, .59]) for the Other cost disparities.

We also tested the hypothesis that people might make decisions that they consider to be fair, regardless of efficiency. This is a reasonable alternative, since we know that humans are sensitive to fairness, for example preferring egalitarian over unequal distributions of income (Dawes et al., 2007). With regards to action planning, it is possible that rather than summing the individual action costs of Self and Other, the decision-maker chooses an action plan to minimize the asymmetry between the individual action costs of the two actors[55]. However, further analyses suggested that fairness was a less likely predictor of decisions made in our experiments, and that in general, the participants prioritized co-efficiency and the minimization of the summed Self and Other costs. This is consistent with the results of Strachan and Török (2020).

The findings from Chapter 3 strengthen previous results from a computational study aimed at formalizing cooperation and competition in a hierarchical model (Kleiman-Weiner et al., 2016). The authors had human participants play coordination games and social dilemma-like games by moving in a grid world, for monetary incentives. Their model defined joint utilities in cooperation

[55] This is also a potential alternative explanation for the door-holding study by Santamaria and Rosenbaum (2011). A door-holder could have focused on both their own and the follower's effort without summing them: the decreasing probability of holding a door open when a follower was further away from the door, rather than closer, could have been due to the door-holder wanting to reduce the follower's individual effort, but not wanting to pay larger extra costs in terms of waiting time. This could be expressed roughly as reasoning that *"The follower is still too far away, by holding the door open for him, I would reduce their effort of going through it, but I would incur a much larger cost in waiting time until he gets here – so perhaps I should not hold the door open."* Such reasoning would consider both person's individual action costs, but not the aggregate shared effort.

as a combination of equally weighted (.5-.5 weights) individual utilities, provided there is no social hierarchy between co-actors. The model captured the observed rates of cooperation and competition well. It is possible that providing environments where one of the co-actors is consistently facing relatively higher costs than the other, asymmetric weights might be necessary for co-efficient decisions, because these higher costs must be minimized to minimize joint costs. However, the randomized design of our task did not provide such consistently asymmetric environments, therefore, co-efficiency could be achieved through equal weighting. Our results were also qualitatively consistent with the findings from Chapter 2, and in line with economic games that found that people sometimes make joint utility-maximizing decisions when monetary rewards are at stake (Colman et al., 2008a).

Chapter 3 raises multiple questions for future research. First, since we deliberately used actions that do not require specific skills, we cannot make any claims about how people would weight the costs of Self and Other when co-actors have largely different competences in performing a given action. It would be interesting to explore if and how actors integrate their partners' relative competence at specific motor tasks into the joint utility of the dyad. One hypothesis is that when experts at a movement (e.g., chefs peeling potatoes) are paired with non-experts, they opt to take on the majority of the costs related to that action, unequally weighting individual utilities in an asymmetric context (as mentioned above). This could also be a way to characterize teaching actions. A comparison of non-teaching and teaching joint actions with competence asymmetries could shine a light on how the higher-level intention to teach a skill to someone (whilst also reaching a shared goal, like peeling the potatoes for dinner) might change the prioritization of co-efficiency. On the one hand, teaching involves communication, which sometimes relies on the use of inefficient actions to inform an interaction partner (see Pezzulo et al., 2019 for a review on sensorimotor communication), and therefore cost-minimization may not be prioritized in teaching, as the benefits of communication through inefficient actions could outweigh the benefits of instrumental efficiency. On the other hand, teachers must have an upper limit on deviating from (co-)efficiency if they want to achieve an instrumental goal while teaching. Research comparing non-teaching and teaching joint actions from the perspective of cost-minimization could elucidate the trade-offs that people make between communication and co-efficiency as a function of higher-level goals. Therefore, such a study could contribute to the exploration of the boundary conditions of co-efficient action planning.

Secondly, boundary conditions may also be researched by addressing the question related to the uncertainty about a partner's cost function. Manipulating the uncertainty about a partner's individual action costs might make people down-weight the importance of those costs in the computation of joint costs, or to ignore them. Similarly, more extreme costs or larger asymmetries between individuals might make people more self-interested or make them focus more on fairness rather than co-efficiency, respectively. Finally, future research should investigate whether people maximize co-efficiency in simultaneous joint actions, too. Results from the co-representation literature (e.g., Schmitz et al., 2017) cause us to speculate that when acting at the same time, people might also account for joint costs.

6.3 Integrating the costs of different types of actions

In Chapter 4 we investigated if the co-efficiency hypothesis also holds in more naturalistic contexts where dyads have to combine different action types to achieve a shared goal. In everyday life, it is often the case that multiple different but complementary sub-tasks make up an action sequence, for example, in cooking (see Wang et al., 2020, for an example of using cooking as a problem for a model of coordination in multi-agent collaboration). Therefore, exploring how people plan such composite joint actions contributes to a better understanding of joint actions beyond the lab.

The study was based on two strands of psychological research, one of which systematically describes how people select actions from among different alternatives (e.g., Feghhi & Rosenbaum, 2019, 2020; Potts, Callahan-Flintoft, et al., 2018; Rosenbaum, 2008, 2012; Rosenbaum et al., 2011).

Rosenbaum and colleagues (2011) argue that a hidden "common currency" is minimized by decisions and in several works, their lab used a psychophysical approach to determine the judged relative costs of actions. The other strand of relevant research contributes to our understanding of how and exactly what people represent about another actor's actions and task constraints when they engage in parallel or (simultaneous or sequential) joint actions with each other (e.g., Schmitz et al., 2017, 2018; Sebanz et al., 2003, 2005; van der Wel & Fu, 2015; Vesper et al., 2013, and the literature on facilitation, see section 6.1 Maximizing co-efficiency in coordination). Notably, there is some evidence that people represent the potential efforts of a co-actor and try to minimize this uncertain effort (Ray et al., 2017). Based on this, we hypothesized that even though in the case of composite joint actions, there would be some uncertainty regarding an interaction partner's exact judged relative costs of the two actions available to them, people would likely be able to represent them. We predicted that people who, in an individual action context, select between two different action types consistently to minimize a common currency, will also take into account a co-actor's action costs in joint action planning. To test this hypothesis, we used a modified object matching task based on Chapter 3. Participants were required to compare the relative costs of tapping and dragging actions in both individual and joint action contexts. As in Chapter 3, we estimated the weights that participants placed on their own and their co-actor's cost disparities.

The findings suggest that a large proportion of people (65%) made decisions between different action types such that they minimized the duration of trials when they acted alone. Their judged relative costs were used to transform tapping costs to the same distance scale as dragging costs, and the functional distances so defined were used as predictors for the participants' decisions in a joint condition. We found that the best-fitting model for cooperative behavior once again included both Self and Other costs disparities, which is consistent with the results in Chapter 3, and with behavior found in the literature on task co-representation. In contrast to Chapter 3, however, the effect of the decision-maker's own action costs was larger than the partner's costs', suggesting that when the integration of action costs was less straightforward due to having to compare and combine different actions with each other, people focused more on minimizing their own costs than their partner's.

These findings reinforce the view that selecting actions from among different alternatives is an apples-and-oranges problem (Rosenbaum & Feghhi, 2019) and it is not evident that all people will always want or manage to minimize the same hidden common currency. How do people compare the costs of joint actions composed of different types of actions? This remains to be explored by future research, in particular the role of simulation, as discussed in section 4.4 Discussion. Previous research seems to support contrary predictions regarding simulation: on the one hand, the representation of a partner's potential efforts benefitted from experience with the task of the partner and subsequently, improved simulation (Ray et al., 2017). On the other hand, Rosenbaum (2012) concluded based on reaction time data that processes other than simulation may be used when people compare composite action sequences (e.g., a two-stage process where people would accept a very low-cost option or reject a high-cost one, but in a second stage, more detailed comparisons might also be made between the potential action plans, if there is no acceptance/rejection in the first stage). Furthermore, future research should combine cognitive effort with physical costs to test if people are able to make co-efficient decisions when some of the costs have no immediately available physical proxies.

6.4 Co-efficiency as potential focal point

Chapter 5 aimed to capture a different aspect of rationality in joint action planning, specifically the potential functions of co-efficient decision-making beyond instrumental benefits. We tested the hypothesis that when interaction partners are not sharing the same space and do not have any information about their partners' actions or an opportunity to communicate, a co-efficient action plan might help coordination by being recognizable as a focal point. Focal points are salient payoff-irrelevant features of strategy pairs that help players converge on the same solution in coordination problems (Lewis, 1969/2008; Schelling, 1960). Salience can be of different kinds, but we focused on Schelling-salience, which refers to the use of a selection rule for coordination, provided that both players recognize it in a given problem (Mehta et al., 1994a).

We hypothesized that co-efficiency could act as a Schelling-salient rule for selecting an action plan that increases the probability of successful coordination. We could phrase the rule as "*Choose the object associated with the minimum possible expected joint movement cost*". Furthermore, co-efficiency could be a payoff-salient focal point, as it is related to action costs by definition. Therefore, we proposed that the minimization of joint action costs would be payoff- and Schelling-salient, which would create a link between Schelling's label-salient focal point concept and coordination on payoff-dominant solutions.

To the best of our knowledge, coordination games have not previously used action-related focal points. Rather, they used perceptual features (e.g., colors), semantic information that are general knowledge (e.g., numbers, notable dates, cities, Mehta et al., 1994a, 1994b) or monetary payoff matrices presented in tables or in spatial arrangements of token-like objects (e.g., Bacharach, 2006; Isoni et al., 2013, 2019). One of the widely researched questions in these latter studies pertains to whether payoff asymmetries between coordinating players help or hinder coordination, and the literature has been inconclusive on this matter. This is relevant to our case of co-efficiency since following a joint-cost minimizing action plan often requires individual inefficiencies from one of the co-actors (i.e., there is often an asymmetry in action costs). This added to our motivation to test the "co-efficiency as focal point" hypothesis.

Using a version of the object matching task in Chapter 3 adapted to online testing, we presented participants with a selection of environments where the expected individual and joint costs of "collecting" objects were manipulated. In two experiments, we asked the participants to match their object choices with a remote partner. Importantly, no feedback was provided about the partner's trial-by-trial decisions to prevent learning. The two experiments differed with regards to the actions required to collect the objects that the participant thought matched what their co-actor was choosing. Experiment 1 used the same three-step procedure as Chapter 3's task, to enable the choice of co-efficient decisions by providing certain information on action costs – anchored to the starting point of the actions. In contrast, in Experiment 2, a one-step procedure provided uncertain information about a partner's starting location, rendering the calculation of joint costs impossible.

We found that the participants' mean proportions of expected coordination success with others were higher than chance in both experiments, and the proportions of object choices that were part of the co-efficient action plans were higher than 50% - this latter being a moderate-sized effect in both experiments. The higher the ratios of co-efficient choices, the likelier the expected coordination success was. Both the probability of coordination success and of co-efficient choices were higher in Experiment 1, where by design, joint costs could be calculated, than in Experiment 2. These effects were complemented by a strong tendency of people to choose the object closer to their cursor when the trial began. However, the parameter estimations suggested that decisions in both experiments were best explained by the combination of Self and Other costs (as defined in Chapter 3) and the object's Distance from Cursor. We concluded that in both experiments, some people seemed to recognize and use co-efficiency as a selection rule for coordination. Although we had hoped that Experiment 2 could serve as a control experiment, participants seemed to base their decisions on costs anchored to the starting locations, which we did not intend to happen. The findings suggest that in case of both real and imagined movement sequences, there was a moderate effect of joint-cost minimization, and co-efficiency may be used by some, but not all people as a focal point.

This study extended the investigation of focal points to the domain of action planning by using the spatial dimensions of stimuli as proxies for action costs which directly relate to payoffs, as opposed to previous work which used them as payoff-irrelevant features. The findings contribute to the line of those studies that did not find a detrimental effect of payoff asymmetries on focal point-based coordination. As mentioned in Chapter 5, the primary open question for future research focuses on the type of reasoning behind the recognition of co-efficiency as a focal point. In theory, both team reasoning (Bacharach, 2006) and reacting to the anticipated actions of an assumed co-efficiency maximizing partner in a cognitive hierarchy (Camerer et al., 2004) could explain this behavior. These two potential underlying processes would explain coordination from different approaches to reasoning: team reasoning being based on an assumption of team agency, whereas cognitive hierarchy theory describes an individualistic mode of reasoning. It is up to future research to explore which type of reasoning might produce decisions based on a co-efficiency selection rule in a coordination game.

6.5 General questions

Before concluding the dissertation, I would like to address a few questions concerning the nature of co-efficiency in general, going beyond particular studies. First, throughout this dissertation I used the terms *rational(ity)* and *co-efficient/-cy* interchangeably. What is the difference between the two? As set out in the theory of infants' teleological reasoning, the rationality principle is a criterion of well-formedness and a productive inferential principle for both the mentalistic and the teleological stances of action understanding (Gergely & Csibra, 2003). The difference between the two stances stems from the representations that the rationality principle applies to. Using the mentalistic stance, adults and children older than 4 attribute goals to others based on representations of the contents of other people's mental states – beliefs about the environment's constraints, desired world states to be reached, and intentions for actions to achieve these desires. In the teleological stance, the three aspects are directly related to reality rather than the contents of mental states: the situational constraints, goal states and actions are held together in a representational schema by the principle of rationality (and because infants are able to represent these, they engage in teleological, not mentalistic reasoning). Applying both stances, inferences can be made about a missing element based on the other two in the same way, for example, actions/intentions can be rationally justified by the combination of the constraints of the environment/beliefs about them and an agent's goal states/desires. In many cases, rational and efficient actions are the same, as they achieve a goal state/desire by the most economical action/intention. However, they differ when the contents represented by mental states do not correspond to reality, for example in cases of false belief or reasoning based on pretense (Gergely & Csibra, 2003). In this dissertation, for simplicity's sake we assumed that people engaging in joint actions have true beliefs about the states of the world, therefore, if their decision-making was rational, it would be also co-efficient. When the possibility of mistaken beliefs and counterfactual realities are addressed, a distinction between rationality and efficiency must be made. However, in the current work we have no basis for making claims about the belief states of co-actors with regards to its match with reality.

Second, the present studies only addressed the part of Santamaria and Rosenbaum's (2011) shared-effort hypothesis that pertains to the observed behavior. We did not test the assumption that sharing beliefs about the odds of potential individual efforts exceeding the potential aggregate efforts was necessary for joint-cost minimization to emerge. Future work might address this question by manipulating the symmetry in the knowledge of co-actors in a joint task. This could be done by inducing false beliefs in the actor who does not make the decision about the action sequence, whilst the decision-maker is aware of this, and has true beliefs about the situation. If sharing beliefs is a precondition for reducing shared effort, then in such an asymmetric context, the effect should disappear. If sharing beliefs about costs is not necessary and people make co-efficient decisions regardless of a co-actor's beliefs about the costs, that would suggest that jointly rational decisions are likely not made for reputational purposes and go beyond physical etiquette, for which the shared-effort hypothesis was proposed as a potential basis (Santamaria & Rosenbaum, 2011).

Third, what is the relationship between co-efficiency and social preferences (e.g., narrow self-interest, altruism, social-welfare, competitive preferences or inequity aversion, Charness & Rabin, 2002)? One might interpret the results we found supporting co-efficient decision-making as not being based on joint costs, but on individuals trying to selfishly minimize their time spent in the lab, without any regard to the dyad's efficiency. According to this reading, co-efficient behavior would ultimately be based on a selfish preference. Certainly, there is an overlap of interests in the experiments that we conducted, since being co-efficient resulted in a shorter total time spent in the lab than when being sub-efficient. This was the case both for the dyad as a unit, and for the individuals making up the dyad. Fixing the total length of the experimental session and communicating this clearly to participants could help tease apart these overlapping interests. If participants still show co-efficient behavior, then it is likely not due to "selfish" lab-time minimization. This is an empirical question, but I would predict that people would still minimize the dyad's time spent on a trial. However, one could also argue that this would only displace the issue to another level (i.e., participants could selfishly minimize the time spent using the touchscreen).

I think the solution lies in making a clear distinction between the implementation of co-efficient actions and the motivations underlying such behaviors, and not considering co-efficiency a type of social preference incompatible with selfishness or altruism. It is possible that different motivational states describable by social preferences like altruism or selfishness result in co-efficiency. In my opinion, decision-makers focusing on joint costs and joint rationality in action planning is a long-term mutualistic strategy for engaging in cooperation. The more often an individual acts co-efficiently and the more often others do the same, the more often they will all benefit from it, regardless of whether co-actors perceive or establish reciprocal relationships between themselves. Therefore, both locally selfish and altruistic behaviors and their underlying social preferences can contribute to this mutualistic, co-efficient strategy, the benefits of which unfold over time. Furthermore, despite unfolding over time, I argue that this does not mean that it is based on direct reciprocity (as in "*I reduce your costs now and you will reduce mine in the next interaction*"), since the non-reciprocal experiment in Chapter 2 found that co-efficient decisions were still likelier than chance when the same one person per dyad made decisions throughout the experiment. In addition, in Chapter 3, the co-efficiency of a partner's previous choice did not improve the predictive power of models that already accounted for the combination of both individual action costs. These results are consistent with the door-holding example in general, too, since that also constitutes a situation where people rarely expect benefits from directly reciprocal behavior. Rather, co-efficient behavior could be an example of indirect reciprocity (Nowak & Sigmund, 2005).

6.6 Conclusions

In this dissertation, I aimed to connect different areas of psychological research by examining the question of rational decision-making in the context of social interactions. Such research has been done in behavioral economics (Colman et al., 2008a, 2008b; Gueye et al., 2020), although in contrast to that investigation, my focus was on the action domain. Rational decisions in the action domain have been a subject of investigation by using monetary rewards, in individual contexts (Trommershäuser et al, 2008). My dissertation complemented this by omitting monetary rewards and manipulating action costs, as well as focusing on joint actions. The closest to my project in these terms were recent studies that used classical games translated into simultaneous sensorimotor tasks (e.g., Braun et al., 2009). We add to that literature by extending the scope of the investigation to sequential joint actions, where co-actors' movements are not coupled with each other. Therefore, in the situations described in this dissertation, there was no direct physical effect of one person's action-related costs on the other person's costs, yet people considered both factors in their decision-making.

I argued that joint actions ought to conform to a principle of rational joint action that states that joint actions should be performed in a way that minimizes the joint costs of the actions sequence, that is, in a co-efficient manner. This would make cooperation more instrumentally efficient and somewhat more predictable for interaction partners. Our findings suggest that the expectation that social interactions unfold rationally is based in actual behavior. In this dissertation, we tested four hypotheses in connection with the principle of rationality and found evidence that when potential costs are easily comparable between individual and joint actions, people plan cooperative actions that minimize the joint costs of a dyad. This becomes more complicated when shared goals require people to perform actions composed of multiple different kinds of actions, and decision-makers will tend to focus more on the minimization of their own costs, as well as estimate their interaction partner's costs based on their own cost function. Finally, under high levels of uncertainty regarding an interaction partner's actions, people might use co-efficiency as a focal point to match their partner's decisions, although this may be just a moderate effect. I argue based on these studies that considering motor decisions in joint action planning as embedded in a mutualistic context may prove fruitful for the exploration of the nature of cooperation. In sum, these findings offer new insight into joint action planning processes from a movement economics perspective and raise intriguing further questions regarding rational choices in cooperation.

References

Axelrod, R., & Hamilton, W. D. (1981). The evolution of cooperation. *Science*, *211*(4489), 1390-1396. doi:10.1126/science.7466396

Bacharach, M. (2006). In: Gold, N., Sugden, R. (Eds.), *Beyond individual choice: teams and frames in game theory*. Princeton University Press.

Bacharach, M., & Bernasconi, M. (1997). The variable frame theory of focal points: An experimental study. *Games and Economic Behavior*, *19*(1), 1-45. doi: 10.1006/game.1997.0546

Bahrami, B., Olsen, K., Bang, D., Roepstorff, A., Rees, G., & Frith, C. (2012). Together, slowly but surely: The role of social interaction and feedback on the build-up of benefit in collective decision-making. *Journal of Experimental Psychology: Human Perception and Performance*, *38*(1), 3. doi:10.1037/a0025708

Bahrami, B., Olsen, K., Latham, P. E., Roepstorff, A., Rees, G., & Frith, C. D. (2010). Optimally interacting minds. *Science, 329(5995),* 1081-1085. doi: 10.1126/science.1185718

Baker, C. L., Jara-Ettinger, J., Saxe, R., & Tenenbaum, J. B. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, *1*(4), 1-10. doi: 10.1038/s41562-017-0064

Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, *113*(3), 329-349. doi: 10.1016/j.cognition.2009.07.005

Bardsley, N., Mehta, J., Starmer, C., & Sugden, R. (2010). Explaining focal points: cognitive hierarchy theory versus team reasoning. *The Economic Journal*, *120*(543), 40-79. doi: 10.1111/j.1468-0297.2009.02304.x

Bates, D., Mächler, M., Bolker, B., Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software, 67*(1), 1-48. doi: 10.18637/jss.v067.i01

Battaglia, P. W., & Schrater, P. R. (2007). Humans trade off viewing time and movement duration to improve visuomotor accuracy in a fast reaching task. *Journal of Neuroscience, 27*(26), 6984-6994. doi: 10.1523/JNEUROSCI.1309-07.2007

Begus, K., Curioni, A., Knoblich, G., & Gergely, G. (2020). Infants understand collaboration: Neural evidence for 9-month-olds 'attribution of shared goals to coordinated joint actions. *Social Neuroscience*, *15*(6), 655-667. doi: 10.1080/17470919.2020.1847730

Blackwell, D., & Girshick, M. A. (1954). *Theory of Games and Statistical Decisions.* Wiley, New York.

Bosch-Domènech, A., & Vriend, N. J. (2013). On the role of non-equilibrium focal points as coordination devices. *Journal of Economic Behavior & Organization*, *94*, 52-67. doi: 10.1016/j.jebo.2013.07.014

Böckler, A., Knoblich, G., & Sebanz, N. (2012). Effects of a coactor's focus of attention on task performance. *Journal of Experimental Psychology: Human Perception and Performance*, *38*(6), 1404. doi: 10.1037/a0027523

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433-436.

Brass, M., Bekkering, H., Wohlschläger, A., & Prinz, W. (2000). Compatibility between observed and executed finger movements: comparing symbolic, spatial, and imitative cues. *Brain and Cognition*, *44*(2), 124-143. doi: 10.1006/brcg.2000.1225

Bratman, M. (1987). *Intention, plans, and practical reason.* Cambridge, MA: Harvard University Press.

Bratman, M. E. (1992). Shared cooperative activity. *The Philosophical Review*, *101*(2), 327-341. doi:10.2307/2185537

Braun, D. A., Ortega, P. A., & Wolpert, D. M. (2009). Nash equilibria in multi-agent motor interactions. *PLoS Computational Biology*, *5*(8), e1000468. doi: 10.1371/journal.pcbi.1000468

Braun, D. A., Ortega, P. A., & Wolpert, D. M. (2011). Motor coordination: when two have to act as one. *Experimental Brain Research*, *211*(3-4), 631-641. doi: 10.1007/s00221-011-2642-y

Brennan, S. E., Chen, X., Dickinson, C. A., Neider, M. B., & Zelinsky, G. J. (2008). Coordinating cognition: The costs and benefits of shared gaze during collaborative search. *Cognition*, *106*(3), 1465-1477. doi: 10.1016/j.cognition.2007.05.012

Burnham, K. P., & Anderson, D. R. (2004). Multimodel inference: understanding AIC and BIC in model selection. *Sociological Methods & Research*, *33*(2), 261-304. doi: 10.1177/0049124104268644

Butterfill, S. (2016). Joint action: A minimal account. In J. Kiverstein (Ed.), *Routledge Handbook of Philosophy of the Social Mind* (pp. 357-369). New York, NY: Routledge.

Camerer, C. F., Ho, T. H., & Chong, J. K. (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, *119*(3), 861-898. doi: 10.1162/0033553041502225

Charness, G., & Rabin, M. (2002). Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, *117*(3), 817-869. doi: 10.1162/003355302760193904

Chennells, M., & Michael, J. (2018). Effort and performance in a cooperative activity are boosted by perception of a partner's effort. *Scientific Reports*, *8*. doi:10.1038/s41598-018-34096-1

Christopoulos, V. N., & Schrater, P. R. (2009). Grasping objects with environmentally induced position uncertainty. *PLoS Computational Biology*, *5*(10), e1000538, doi:10.1371/journal.pcbi.1000538

Cohen, R. G., Biddle, J. C., & Rosenbaum, D. A. (2010). Manual obstacle avoidance takes into account visual uncertainty, motor noise, and biomechanical costs. *Experimental Brain Research*, *201*(3), 587-592. doi: 10.1007/s00221-009-2042-8

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.

Colman, A. M. (1997). Salience and focusing in pure coordination games. *Journal of Economic Methodology*, *4*(1), 61-81. doi: 10.1080/13501789700000004

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

Colman, A. M., Pulford, B. D., & Rose, J. (2008b). Team reasoning and collective rationality: Piercing the veil of obviousness. *Acta Psychologica, 128,* 409-412. doi: 10.1016/j.actpsy.2008.04.001

Constable, M. D., Bayliss, A. P., Tipper, S. P., Spaniol, A. P., Pratt, J., & Welsh, T. N. (2016). Ownership status influences the degree of joint facilitatory behavior. *Psychological Science, 27*(10), 1371-1378. doi:10.1177/0956797616661544

Cooper, R. W., DeJong, D. V., Forsythe, R., & Ross, T. W. (1990). Selection criteria in coordination games: Some experimental results. *The American Economic Review*, *80*(1), 218-233.

Crawford, V. P., Gneezy, U., & Rottenstreich, Y. (2008). The power of focal points is limited: Even minute payoff asymmetry may yield large coordination failures. *American Economic Review*, *98*(4), 1443-58. doi: 10.1257/aer.98.4.1443

Csibra, G. (2017). Cognitive science: modelling theory of mind. *Nature Human Behaviour*, *1*(4), 1-1. doi: 10.1038/s41562-017-0066

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. doi: 10.1016/S0364-0213(02)00112-X

Csibra, G., & Gergely, G. (2007). 'Obsessed with goals': Functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychologica*, *124*(1), 60-78. doi: 10.1016/j.actpsy.2006.09.007

Curioni, A., Voinov, P., Allritz, M., Call, J., & Knoblich, G. K. (2020). Crazy for you! Understanding Utility in Joint Actions. In S. Denison., M. Mack, Y. Xu, & B.C. Armstrong (Eds.), *Proceedings of the 42nd Annual Conference of the Cognitive Science Society* (pp. 3268-3274.). Cognitive Science Society.

Cuthill, I. C., & Houston, A. I. (1997). Managing time and energy. In J. R. Krebs & N. B. Davies (Eds.), *Behavioural ecology: An evolutionary approach* (pp. 97–120). Oxford, England: Blackwell Science.

Davis, M. H. (1980). A multidimensional approach to individual differences in empathy. *JSAS Catalog of Selected Documents in Psychology, 10*, 85.

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

Della Gatta, F., Garbarini, F., Rabuffetti, M., Viganò, L., Butterfill, S.A., Sinigaglia, C. (2017). Drawn together: when motor representations ground joint actions. *Cognition 165*, 53–60. doi: 10.1016/j.cognition.2017.04.008.

Dennett, D. C. (1987). *The intentional stance*. MIT press.

Denwood, M. J. (2016). runjags: An R package providing interface utilities, model templates, parallel computing methods and additional distributions for MCMC models in JAGS. *Journal of Statistical Software, 71(9)*, 1-25. doi:10.18637/jss.v071.i09

Dötsch, D., & Schubö, A. (2015). Social categorization and cooperation in motor joint action: evidence for a joint end-state comfort. *Experimental Brain Research*, *233*(8), 2323-2334. doi: 10.1007/s00221-015-4301-1

Dunn, T. L., Lutes, D. J., & Risko, E. F. (2016). Metacognitive evaluation in the avoidance of demand. *Journal of Experimental Psychology: Human Perception and Performance*, *42*(9), 1372. doi: 10.1037/xhp0000236

Elekes, F., Bródy, G., Halász, E., & Király, I. (2016). Enhanced encoding of the co-actor's target stimuli during a shared non-motor task. *Quarterly Journal of Experimental Psychology*, *69*(12), 2376-2389. doi: 10.1080/17470218.2015.1120332

Elliott, D., Hansen, S., Mendoza, J., & Tremblay, L. (2004). Learning to optimize speed, accuracy, and energy expenditure: a framework for understanding speed-accuracy relations in goal-directed aiming. *Journal of Motor Behavior*, *36*(3), 339-351. doi: 10.3200/JMBR.36.3.339-351

Engemann, D. A., Bzdok, D., Eickhoff, S. B., Vogeley, K., & Schilbach, L. (2012). Games people play—toward an enactive view of cooperation in social neuroscience. *Frontiers in Human Neuroscience*, *6*, 148. doi: 10.3389/fnhum.2012.00148

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*(6870), 429-433. doi: 10.1038/415429[a]

Faillo, M., Smerilli, A., & Sugden, R. (2017). Bounded best-response and collective-optimality reasoning in coordination games. *Journal of Economic Behavior & Organization*, *140*, 317-335. doi: 10.1016/j.jebo.2017.05.015

Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*, 175-191. doi:10.3758/BF03193146

Fawcett, C., & Gredebäck, G. (2013). Infants use social context to bind actions into a collaborative sequence. *Developmental Science*, *16*(6), 841–849. doi: 10.1111/desc.12074

Fawcett, C., & Gredebäck, G. (2015). Eighteen-month-olds, but not 14-month-olds, use social context to bind action sequences. *Infancy*, *20*(1), 115–125. doi: 10.1111/infa.12062

Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters*, *27*(8), 861-874. doi:10.1016/j.patrec.2005.10.010

Feghhi, I., & Rosenbaum, D. A. (2019). Judging the subjective difficulty of different kinds of tasks. *Journal of Experimental Psychology: Human Perception and Performance*, *45*(8), 983. doi: 10.1037/xhp0000653

Feghhi, I., & Rosenbaum, D. A. (2020). Effort avoidance is not simply error avoidance. *Psychological Research*. doi: 10.1007/s00426-020-01331-2

Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, *114*(3), 817-868. doi: 10.1162/003355399556151

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. doi: 10.1037/h0055392

Fournier, L. R., Coder, E., Kogan, C., Raghunath, N., Taddese, E., & Rosenbaum, D. A. (2019). Which task will we choose first? Precrastination and cognitive load in task ordering. *Attention, Perception, & Psychophysics*, *81*(2), 489-503. doi: 10.3758/s13414-018-1633-5

Fournier, L. R., Stubblefield, A. M., Dyre, B. P., & Rosenbaum, D. A. (2019). Starting or finishing sooner? Sequencing preferences in object transfer tasks. *Psychological Research*, *83*(8), 1674-1684. doi: 10.1007/s00426-018-1022-7

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

Gan, K. C., & Hoffmann, E. R. (1988). Geometrical conditions for ballistic and visually controlled movements. *Ergonomics*, *31*(5), 829-839. doi: 10.1080/00140138808966724

Gärling, T., & Gärling, E. (1988). Distance minimization in downtown pedestrian shopping. *Environment and Planning A*, *20*(4), 547-554. doi:10.1068/a200547

Gelman, A., Hwang, J., & Vehtari, A. (2014). Understanding predictive information criteria for Bayesian models. *Statistics and Computing*, *24*(6), 997-1016. doi: 10.1007/s11222-013-9416-2

Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science, 7*, 457-511.

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

Gergely, G., Nádasdy, Z., Csibra, G., & Bíró, S. (1995). Taking the intentional stance at 12 months of age. *Cognition*, *56*(2), 165-193. doi: 10.1016/0010-0277(95)00661-H

Gilbert, M. (1987). Modelling collective belief. *Synthese*, *73*(1), 185-204. doi:10.1007/BF00485446

Gilbert, M. (2006). Rationality in collective action. *Philosophy of the Social Sciences*, *36*(1), 3-17. Doi: 10.1177/0048393105284167

Gonzalez, D. A., Studenka, B. E., Glazebrook, C. M., & Lyons, J. L. (2011). Extending end-state comfort effect: Do we consider the beginning state comfort of another? *Acta Psychologica*, *136*(3), 347-353. doi:10.1016/j.actpsy.2010.12.009

Grau-Moya, J., Hez, E., Pezzulo, G., & Braun, D. A. (2013). The effect of model uncertainty on cooperation in sensorimotor interactions. *Journal of The Royal Society Interface*, *10*(87), 20130554. doi: 10.1098/rsif.2013.0554

Gredebäck, G., & Melinder, A. (2010). Infants 'understanding of everyday social interactions: A dual process account. *Cognition*, *114*(2), 197-206. doi: 10.1016/j.cognition.2009.09.004

Gueye, M., Quérou, N., & Soubeyran, R. (2020). Social preferences and coordination: An experiment. *Journal of Economic Behavior & Organization*, *173*, 26-54. doi: 10.1016/j.jebo.2020.02.017

Hampton, J. (1998). Rationality, practical. In *The Routledge Encyclopedia of Philosophy*. Taylor and Francis. Retrieved 5 Feb. 2021, from https://www.rep.routledge.com/articles/thematic/rationality-practical/v-1. doi:10.4324/9780415249126-W035-1

Harris, C. M., & Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, *394*(6695), 780. doi: 10.1038/29528.

Harrison, X. A., Donaldson, L., Correa-Cano, M. E., Evans, J., Fisher, D. N., Goodwin, C. E. D., Robinson, B. S., Hodgson, D. J., Inger, R. (2018). A brief introduction to mixed effects modelling and multi-model inference in ecology. *PeerJ* 6:e4794 doi: 10.7717/peerj.4794

Harsanyi, J., & Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, Mass.

Heider, F., & Simmel, M. (1944). An experimental study of apparent behavior. *The American Journal of Psychology*, *57*(2), 243-259. doi: 10.2307/1416950

Henderson, A. M. E. , Wang, Y. , Matz, L. E. , & Woodward, A. L. (2013). Active experience shapes 10-month-old infants 'understanding of collaborative goals. *Infancy*, *18*(1), 10–39. doi: 10.1111/j.1532-7078.2012.00126.x

Hommel, B., Müsseler, J., Aschersleben, G., & Prinz, W. (2001). The theory of event coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences*, *24*(5), 849. doi: 10.1017/s0140525x01000103.

Hoyt, D. F., & Taylor, C. R. (1981). Gait and the energetics of locomotion in horses. *Nature, 292(5820), 239–240.* doi:10.1038/292239a0

Hull, C. L. (1943). *Principles of behavior: an introduction to behavior theory.* New York: Appleton-Century.

Hurley, S. (2005). Social heuristics that make us smarter. *Philosophical Psychology*, *18*(5), 585-612. doi:10.1080/09515080500264214

Isoni, A., Poulsen, A., Sugden, R., & Tsutsui, K. (2013). Focal points in tacit bargaining problems: Experimental evidence. *European Economic Review*, *59*, 167-188. Doi: 10.1016/j.euroecorev.2012.12.005

Isoni, A., Poulsen, A., Sugden, R., & Tsutsui, K. (2019). Focal points and payoff information in tacit bargaining. *Games and Economic Behavior*, *114*, 193-214. doi: 10.1016/j.geb.2019.01.008

Isoni, A., Sugden, R., & Zheng, J. (2018). *The Pizza Night Game: Efficiency, Conflict and Inequality in Tacit Bargaining Games with Focal Points* (No. 18-01). School of Economics, University of East Anglia, Norwich, UK.

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. doi: 10.1016/j.tics.2016.05.011

Jara-Ettinger, J., Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2015). Children's understanding of the costs and rewards underlying rational action. *Cognition*, *140*, 14-23. doi: 10.1016/j.cognition.2015.03.006

Jara-Ettinger, J., Schulz, L. E., & Tenenbaum, J. B. (2020). The naive utility calculus as a unified, quantitative framework for action understanding. *Cognitive Psychology*, *123*, 101334. doi: 10.1016/j.cogpsych.2020.101334

Jarvstad, A., Hahn, U., Rushton, S. K., & Warren, P. A. (2013). Perceptuo-motor, cognitive, and description-based decision-making seem equally good. *Proceedings of the National Academy of Sciences*, *110*(40), 16271-16276. doi: 10.1073/pnas.1300239110

JASP Team (2018). JASP (Version 0.8.5) [Computer software].

Juslin, P., & Olsson, H. (1997). Thurstonian and Brunswikian origins of uncertainty in judgment: A sampling model of confidence in sensory discrimination. *Psychological Review, 104(2),* 344–366. doi: 10.1037/0033-295X.104.2.344

Kerby, D. S. (2014). The simple difference formula: An approach to teaching nonparametric correlation. *Comprehensive Psychology, 3,* 11-IT. doi: 10.2466/11.IT.3.1

Király, I., Jovanovic, B., Prinz, W., Aschersleben, G., & Gergely, G. (2003). The early origins of goal attribution in infancy. *Consciousness and Cognition*, *12*(4), 752-769. doi: 10.1016/S1053-8100(03)00084-9

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.

Knoblauch, K., & Maloney, L. T. (2012). *Modeling psychophysical data in R* (Vol. 32). Springer Science & Business Media. doi: 10.1007/978-1-4614-4475-6

Knoblich, G., Butterfill, S., & Sebanz, N. (2011). Psychological research on joint action: theory and data. In *Psychology of Learning and Motivation* (Vol. 54, pp. 59-101). Academic Press.

Kool, W., McGuire, J. T., Rosen, Z. B., & Botvinick, M. M. (2010). Decision making and the avoidance of cognitive demand. *Journal of Experimental Psychology: General, 139*, 665–682. doi: 10.1037/a0020198

Körding, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences*, *10*(7), 319-326. doi: 10.1016/j.tics.2006.05.003

Kourtis, D., Knoblich, G., Woźniak, M., & Sebanz, N. (2014). Attention allocation and task representation during joint action planning. *Journal of Cognitive Neuroscience*, *26*(10), 2275-2286. doi: 10.1162/jocn_a_00634

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. doi: 10.1080/17470919.2012.694823

Kourtis, D., Woźniak, M., Sebanz, N., & Knoblich, G. (2019). Evidence for we-representations during joint action planning. *Neuropsychologia*, *131*, 73-83. doi: 10.1016/j.neuropsychologia.2019.05.029

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

Kubrick, S. (Director). (1968). *2001: A Space Odyssey* [Film]. Metro-Goldwyn-Mayer.

Le Bars, S., Bourgeois-Gironde, S., Wyart, V., Sari, I., Pacherie, E., & Chambon, V. (2020 March 24). Motor coordination and strategic cooperation in joint action. Preprint retrieved from https://psyarxiv.com/xbm34/download

Lee, M. D., & Wagenmakers, E. J. (2014). *Bayesian cognitive modeling: A practical course*. Cambridge university press. doi: 10.1017/CBO9781139087759

Lewis, D. (1969/2008). *Convention: A philosophical study*. (2nd ed.) John Wiley & Sons.

Liu, S., Ullman, T. D., Tenenbaum, J. B., & Spelke, E. S. (2017). Ten-month-old infants infer the value of goals from the costs of actions. *Science, 358(6366),* 1038-1041. doi: 10.1126/science.aag2132

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. doi: 10.1007/s00221-006-0426-6

Maloney, L. T., Trommershäuser, J., & Landy, M. S. (2007). Questions without words: A comparison between decision making under risk and movement planning under risk. In W. D. Gray (Ed.), *Series on cognitive models and architectures. Integrated models of cognitive systems* (p. 297–313). Oxford University Press.

Maloney, L. T., & Zhang, H. (2010). Decision-theoretic models of visual perception and action. *Vision Research*, *50*(23), 2362-2374. doi: 10.1016/j.visres.2010.09.031

Mascaro, O. & Csibra, G. (2014). Fourteen-month-old infants compute the efficiency of joint actions. Paper presented at the XIX. Biennial International Conference on Infant Studies, 5 July 2014, Berlin, Germany.

McElreath, R. (2020). *Statistical rethinking: A Bayesian course with examples in* R *and Stan*. CRC press.

McLaughlin, R. L., & Montgomerie, R. D. (1990). Flight speeds of parent birds feeding nestlings: maximization of foraging efficiency or food delivery rate?. *Canadian Journal of Zoology*, *68*(11), 2269-2274. doi: 10.1139/z90-316

Mehta, J., Starmer, C., & Sugden, R. (1994a). The nature of salience: An experimental investigation of pure coordination games. *The American Economic Review*, *84*(3), 658-673.

Mehta, J., Starmer, C., & Sugden, R. (1994b). Focal points in pure coordination games: An experimental investigation. *Theory and Decision*, *36*(2), 163-185. Doi: 10.1007/BF01079211

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. doi: 10.1007/s00221-012-3398-8

Meyer, M., van der Wel, R. P., & Hunnius, S. (2016). Planning my actions to accommodate yours: joint action development during early childhood. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *371*(1693), 20150371. doi: 10.1098/rstb.2015.0371

Misyak, J. B., & Chater, N. (2014). Virtual bargaining: a theory of social decision-making. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *369*(1655), 20130487. doi: 10.1098/rstb.2013.0487

Misyak, J. B., Melkonyan, T., Zeitoun, H., & Chater, N. (2014). Unwritten rules: virtual bargaining underpins social interaction, culture, and society. *Trends in Cognitive Sciences*, *18*(10), 512-519. doi: 10.1016/j.tics.2014.05.010

Nakagawa, S., & Schielzeth, H. (2013). A general and simple method for obtaining R2 from generalized linear mixed-effects models. *Methods in Ecology and Evolution*, *4*(2), 133-142. doi: 10.1111/j.2041-210x.2012.00261.x

Nash, J. (1951). Non-cooperative games. *Annals of Mathematics*, 286-295. doi: 10.2307/1969529

Nash, J. (1953). Two-person cooperative games. *Econometrica: Journal of the Econometric Society*, 128-140. doi: 10.2307/1906951

Neumann, J. von, & Morgenstern, O. (1944). *Theory of Games and Economic Behavior*, Princeton, NJ: Princeton University Press, 1953.

Nowak, M. A., & Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature*, *437*(7063), 1291-1298. doi: 10.1038/nature04131

Parravano, M., & Poulsen, O. (2015). Stake size and the power of focal points in coordination games: Experimental evidence. *Games and Economic Behavior*, *94*, 191-199. doi: 10.1016/j.geb.2015.05.001

Pelli, D. G. (1997) The VideoToolbox software for visual psychophysics: Transforming numbers into movies, *Spatial Vision 10*:437-442.

Pezzulo, G., Donnarumma, F., Dindo, H., D'Ausilio, A., Konvalinka, I., & Castelfranchi, C. (2019). The body talks: Sensorimotor communication and its brain and kinematic signatures. *Physics of Life Reviews*, *28*, 1-21. doi: 10.1016/j.plrev.2018.06.014

Philipp, A. M., & Prinz, W. (2010). Evidence for a role of the responding agent in the joint compatibility effect. *Quarterly Journal of Experimental Psychology*, *63*(11), 2159-2171. doi: 10.1080/17470211003802426

Plummer, M. (2003, March). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd international workshop on distributed statistical computing* (Vol. 124, No. 125.10). doi:10.1.1.13.3406

Potts, C. A., Callahan-Flintoft, C., & Rosenbaum, D. A. (2018). How do reaching and walking costs affect movement path selection?. *Experimental Brain Research*, *236*(10), 2727-2737. doi: 10.1007/s00221-018-5327-y

Potts, C. A., Pastel, S., & Rosenbaum, D. A. (2018). How are cognitive and physical difficulty compared?. *Attention, Perception, & Psychophysics*, *80*(2), 500-511. doi: 10.3758/s13414-017-1434-2

Poulsen, A., & Sonntag, A. (2019). *Focality is Intuitive-Experimental Evidence on the Effects of Time Pressure in Coordination Games* (No. 19-01). School of Economics, University of East Anglia, Norwich, UK.

Raghunath, N., Fournier, L. R., & Kogan, C. (2020). Precrastination and individual differences in working memory capacity. *Psychological Research*, 1-16. doi: 10.1007/s00426-020-01373-6

Rand, D. G., Tarnita, C. E., Ohtsuki, H., & Nowak, M. A. (2013). Evolution of fairness in the one-shot anonymous Ultimatum Game. *Proceedings of the National Academy of Sciences*, *110*(7), 2581-2586. doi: 10.1073/pnas.1214167110

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. doi: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. doi:10.1080/00222895.2011.592871

Ray, M., & Welsh, T. N. (2018). Multiple frames of reference are used during the selection and planning of a sequential joint action. *Frontiers in Psychology*, *9*, 542. doi: 10.3389/fpsyg.2018.00542

Roberts, M. E., & Goldstone, R. L. (2011). Adaptive group coordination and role differentiation. *PLoS One, 6(7)*, e22377. doi: 10.1371/journal.pone.0022377

Rosenbaum, D. A. (2008). Reaching while walking: reaching distance costs more than walking distance. *Psychonomic Bulletin & Review*, *15*(6), 1100-1104. doi: 10.3758/PBR.15.6.1100

Rosenbaum, D. A. (2012). The tiger on your tail: Choosing between temporally extended behaviors. *Psychological Science*, *23*(8), 855-860. doi: 10.1177/0956797612440459

Rosenbaum, D. A., Brach, M., & Semenov, A. (2011). Behavioral ecology meets motor behavior: Choosing between walking and reaching paths. *Journal of Motor Behavior*, *43*(2), 131-136. doi: 10.1080/00222895.2010.548423

Rosenbaum, D. A., & Bui, B. V. (2019). Does task sustainability provide a unified measure of subjective task difficulty?. *Psychonomic Bulletin & Review*, *26*(6), 1980-1987. doi: 10.3758/s13423-019-01631-8.

Rosenbaum, D. A., Chapman, K. M., Coelho, C. J., Gong, L., & Studenka, B. E. (2013). Choosing actions. *Frontiers in Psychology*, *4*, 273. doi: 10.3389/fpsyg.2013.00273

Rosenbaum, D. A., & Feghhi, I. (2019). The time for action is at hand. *Attention, Perception, & Psychophysics*, *81*(7), 2123-2138. doi: 10.3758/s13414-018-01647-7

Rosenbaum, D. A., & Gaydos, M. J. (2008). A method for obtaining psychophysical estimates of movement costs. *Journal of Motor Behavior*, *40*(1), 11-17. Doi: 10.3200/JMBR.40.1.11-17

Rosenbaum, D. A., Gong, L., & Potts, C. A. (2014). Pre-crastination: Hastening subgoal completion at the expense of extra physical effort. *Psychological Science*, *25*(7), 1487-1496. doi: 10.1177/0956797614532657

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* (p. 321–342). Lawrence Erlbaum Associates, Inc.

Rosenbaum, D. A., van Heugten, C. M., & Caldwell, G. E. (1996). From cognition to biomechanics and back: The end-state comfort effect and the middle-is-faster effect. *Acta Psychologica*, *94*(1), 59-85. doi: 10.1016/0001-6918(95)00062-3

Sacheli, L. M., Arcangeli, E., & Paulesu, E. (2018). Evidence for a dyadic motor plan in joint action. *Scientific Reports*, *8*(1), 1-12. doi: 10.1038/s41598-018-23275-9

Sacheli, L. M., Candidi, M., Pavone, E. F., Tidoni, E., & Aglioti, S. M. (2012). And yet they act together: interpersonal perception modulates visuo-motor interference and mutual adjustments during a joint-grasping task. *PloS One*, *7*(11), e50223.

Santamaria, J. P., & Rosenbaum, D. A. (2011). Etiquette and effort holding doors for others. *Psychological Science*, *22*(5*),* 584-588. doi:10.1177/0956797611406444

Scharoun, S., Mintz, A., Glazebrook, C. M., Roy, E., & Gonzalez, D. (2017). How Working With a Helpful Versus Less Helpful Confederate Influences Joint-Action in a Pegboard Task. *Journal of Motor Behavior*, 1-10. doi: 10.1080/00222895.2016.1250715

Scharoun, S. M., Scanlan, K. A., & Bryden, P. J. (2016). Hand and grasp selection in a preferential reaching task: the effects of object location, orientation, and task intention. *Frontiers in Psychology*, *7*.

Schelling, T. C. (1960). *The strategy of conflict*. Harvard University Press.

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. doi: 10.1037/xhp0000403

Schmitz, L., Vesper, C., Sebanz, N., & Knoblich, G. (2018). Co-actors represent the order of each other's actions. *Cognition*, *181*, 65-79. doi: 10.1016/j.cognition.2018.08.008

Scott, R. M., & Baillargeon, R. (2013). Do infants really expect agents to act efficiently? A critical test of the rationality principle. *Psychological Science*, *24*(4), 466-474. doi: 10.1177/0956797612457395

Sebanz, N., Bekkering, H., & Knoblich, G. (2006). Joint action: bodies and minds moving together. *Trends in Cognitive Sciences*, *10*(2), 70-76. doi: 10.1016/j.tics.2005.12.009

Sebanz, N., & Knoblich, G. (2021). Progress in Joint-Action Research. *Current Directions in Psychological Science*, 0963721420984425. doi: 10.1177/0963721420984425

Sebanz, N., Knoblich, G., & Prinz, W. (2003). Representing others' actions: just like one's own?. *Cognition*, *88*(3), B11-B21. doi: 10.1016/s0010-0277(03)00043-x

Sebanz, N., Knoblich, G., & Prinz, W. (2005). How two share a task: corepresenting stimulus-response mappings. *Journal of Experimental Psychology: Human Perception and Performance*, *31*(6), 1234. doi: 10.1037/0096-1523.31.6.1234

Sebanz, N., Knoblich, G., Prinz, W., & Wascher, E. (2006). Twin peaks: An ERP study of action planning and control in coacting individuals. *Journal of Cognitive Neuroscience*, *18*(5), 859-870. doi: 10.1162/jocn.2006.18.5.859

Singmann, H., & Kellen, D. (2019). An Introduction to Mixed Models for Experimental Psychology. In D. H. Spieler & E. Schumacher (Eds.), *New Methods in Cognitive Psychology* (pp. 4–31). Psychology Press.

Skerry, A. E., Carey, S. E., & Spelke, E. S. (2013). First-person action experience reveals sensitivity to action efficiency in prereaching infants. *Proceedings of the National Academy of Sciences*,*110*(46), 18728-18733. doi: 10.1073/pnas.1312322110

Sorkin, R. D., Hays, C. J., & West, R. (2001). Signal-detection analysis of group decision making. *Psychological Review*, *108*(1), 183. doi: 10.1037/0033-295x.108.1.183

Southgate, V., Johnson, M. H., & Csibra, G. (2008). Infants attribute goals even to biomechanically impossible actions. *Cognition*, *107*(3), 1059-1069. doi: 10.1016/j.cognition.2007.10.002

Sparrow, W. A., & Newell, K. M. (1998). Metabolic energy expenditure and the regulation of movement economy. *Psychonomic Bulletin & Review*, *5*(2), 173-196. doi: 10.3758/BF03212943

Stahl, D. O., & Wilson, P. W. (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, *10*(1), 218-254. doi: 10.1006/game.1995.1031

Strachan, J. W., & Török, G. (2020). Efficiency is prioritised over fairness when distributing joint actions. *Acta Psychologica*, *210*, 103158. doi: 10.1016/j.actpsy.2020.103158

Stueber, K. (2010). *Rediscovering empathy: Agency, folk psychology, and the human sciences*. MIT Press.

Sugden, R. (1993). Thinking as a team: Towards an explanation of nonselfish behavior. *Social Philosophy and Policy*, *10*(1), 69-89. doi: 10.1017/S0265052500004027

Sugden, R. (1995). A theory of focal points. *The Economic Journal*, *105*(430), 533-550. doi: 10.2307/2235016

Sugden, R. (2000). Team preferences. *Economics & Philosophy*, *16*(2), 175-204. doi:10.1017/S0266267100000213

Sugden, R. (2003). The logic of team reasoning. *Philosophical Explorations*, *6*(3), 165-181. doi:10.1080/10002003098538748

Summerfield, C., & Tsetsos, K. (2020). Rationality and Efficiency in Human Decision-making. In D., Poeppel, G. R., Mangun, & M. S., Gazzaniga (Eds.), *The Cognitive Neurosciences* (pp. 427-438). Cambridge, MA: MIT Press.

Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature Neuroscience*, *7*(9), 907-915. doi: 10.1038/nn1309

Tolman, E. C. (1932). *Purposive behavior in animals and men.* New York: Century.

Tomasello, M. (2009). *Why we cooperate*. MIT Press.

Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences, 28(5)*, 675-691. doi: 10.1017/S0140525X05000129

Tomasello, M., Kruger, A. C., & Ratner, H. H. (1993). Cultural learning. *Behavioral and Brain Sciences*, *16*(3), 495-511. doi: 10.1017/S0140525X0003123X

Török, G., Pomiechowska, B., Csibra, G., & Sebanz, N. (2019). Rationality in joint action: maximizing coefficiency in coordination. *Psychological Science*, *30*(6), 930-941. doi: 10.1177/0956797619842550

Trivers, R. L. (1971). The evolution of reciprocal altruism. *The Quarterly Review of Biology*, *46*(1), 35-57. doi: 10.1086/406755

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003a). Statistical decision theory and trade-offs in the control of motor response. *Spatial Vision*,*16*(3), 255-275. doi: 10.1163/156856803322467527

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2003b). Statistical decision theory and the selection of rapid, goal-directed movements. *Journal of the Optical Society of America A, 20(7),* 1419–1433.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in Cognitive Sciences*, *12*(8), 291-297. doi: 10.1016/j.tics.2008.04.010

Tsai, C. C., Kuo, W. J., Hung, D. L., & Tzeng, O. J. L. (2008). Action co-representation is tuned to other humans. *Journal of Cognitive Neuroscience*, *20*(11), 2015–2024. doi: 10.1162/jocn.2008.20144

Tsai, C. C., Kuo, W. J., Jing, J. T., Hung, D. L., & Tzeng, O. J. L. (2006). A common coding framework in self-other interaction: Evidence from joint action task. *Experimental Brain Research*, *175*(2), 353–362. Doi: 10.1007/s00221-006-0557-9

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124-1131.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*(4481), 453-458. doi: 10.1126/science.7455683

van der Wel, R. P., & Fu, E. (2015). Entrainment and task co-representation effects for discrete and continuous action sequences. *Psychonomic Bulletin & Review, 22(6)*, 1685–1691. Doi: 10.3758/s13423-015-0831-6

van Elten, J., & Penczynski, S. P. (2020). Coordination games with asymmetric payoffs: An experimental study with intra-group communication. *Journal of Economic Behavior & Organization*, *169*, 158-188. doi: 10.1016/j.jebo.2019.11.006

Van Huyck, J. B., Battalio, R. C., & Beil, R. O. (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *The American Economic Review*, *80*(1), 234-248.

Vehtari, A., Gabry, J., Magnusson, M., Yao, Y., Bürkner, P., Paananen, T., & Gelman, A. (2020). "loo: Efficient leave-one-out cross-validation and WAIC for Bayesian models." R package version 2.4.1, https://mc-stan.org/loo/

Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, *27*(5), 1413-1432. doi: 10.1007/s11222-016-9696-4

Vesper, C., Abramova, E., Bütepage, J., Ciardo, F., Crossey, B., Effenberg, A., ... & Schmitz, L. (2017). Joint action: mental representations, shared information and general mechanisms for coordinating with others. *Frontiers in Psychology*, *7*, 2039. doi: 10.3389/fpsyg.2016.02039

Vesper, C., Butterfill, S., Knoblich, G., & Sebanz, N. (2010). A minimal architecture for joint action. *Neural Networks*, *23*(8-9), 998-1003. doi: 10.1016/j.neunet.2010.06.002

Vesper, C., van der Wel, R. P., 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. doi: 10.1037/a0028066

Voinov, P. V., Sebanz, N., & Knoblich, G. (2019). Collective benefit in joint perceptual judgments: Partial roles of shared environments, meta-cognition, and feedback. *Cognition, 189,* 116–130. doi: 10.1016/j.cognition.2019.03.016

Wagenmakers, E. J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, *11*(1), 192-196. doi: 10.3758/BF03206482

Wahn, B., Kingstone, A., & König, P. (2017). Two trackers are better than one: Information about the co-actor's actions and performance scores contribute to the collective benefit in a joint visuospatial task. *Frontiers in Psychology*, *8*, 669. doi: 0.3389/fpsyg.2017.00669

Wallace, R. J. (2020). Practical Reason. In *The Stanford Encyclopedia of Philosophy*. Ed. Edward N. Zalta, Retrieved 5 Feb. 2021, from https://plato.stanford.edu/archives/spr2020/entries/practical-reason/.

Wang, R. E., Wu, S. A., Evans, J. A., Tenenbaum, J. B., Parkes, D. C., & Kleiman-Weiner, M. (2020). Too many cooks: Coordinating multi-agent collaboration through inverse planning. *arXiv preprint arXiv:2003.11778*.

Watanabe, S., & Opper, M. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning research*, *11*(12).

Welham, C. V., & Ydenberg, R. C. (1993). Efficiency-Maximizing Flight Speeds in Parent Black Terns. *Ecology*, *74*(6), 1893-1901. doi: 10.2307/1939946

Welsh, T. N., Higgins, L., Ray, M., & Weeks, D. J. (2007). Seeing vs. believing: Is believing sufficient to activate the processes of response co-representation?. *Human Movement Science*, *26*(6), 853-866. doi: 10.1016/j.humov.2007.06.003

Wenke, D., Atmaca, S., Holländer, A., Liepelt, R., Baess, P., & Prinz, W. (2011). What is shared in joint action? Issues of co-representation, response conflict, and agent identification. *Review of Philosophy and Psychology*, *2*(2), 147-172. doi: 10.1007/s13164-011-0057-0

Wolf, T., Sebanz, N., & Knoblich, G. (2020). Adaptation to unstable coordination patterns in individual and joint actions. *PloS One*, *15*(5), e0232667. doi: 10.1371/journal.pone.0232667

Wolpert, D. M., & Landy, M. S. (2012). Motor control is decision-making. *Current Opinion in Neurobiology*, *22*(6), 996-1003. doi:10.1016/j.conb.2012.05.003

Wu, S. W., Delgado, M. R., & Maloney, L. T. (2009). Economic decision-making compared with an equivalent motor task. *Proceedings of the National Academy of Sciences*, *106*(15), 6088-6093. doi: 10.1073/pnas.0900102106

Appendix A – Additional Results to Chapter 4

A.1 Individual No Choice Condition

Table A.1: Results of the mixed-effects linear regression model predicting movement durations by Action Type (Tapping vs. Dragging) in the Individual No Choice condition. Raw estimates on a log scale are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Movement Duration ~ Action Type + (Action Type | participant)			
Predictors	*Estimates*	*std. Error*	*z*	*p*
Intercept	-0.232 (-0.312 – -0.152)	0.041	-5.694	**<0.001**
Action Type[Tap]	0.066 (-0.029 – 0.162)	0.049	1.362	0.173
Random Effects				
σ^2	0.49			
$\tau_{00\ subj}$	0.04			
$\tau_{11\ subj.costChoiceTap}$	0.05			
$\rho_{01\ subj}$	-0.56			
ICC	0.08			
N_{subj}	40			
Observations	1440			
Marginal R^2 / Conditional R^2	0.002 / 0.083			
AIC	1422.620			

Table A.2: Results of the mixed-effects linear regression model predicting movement durations by Action Type (Tapping vs. Dragging) * log(Path) * log(Step Number) in the Individual No Choice condition. Raw estimates on a log scale are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations. The reference level for Action Type is *Dragging*.

	Movement Duration ~ Action Type X log(Path) X log(Step Number) + (Action Type X log(Path) | participant)			
Predictors	*Estimates*	*std. Error*	*z*	*p*
Intercept	-5.455 (-7.601 – -3.310)	1.095	-4.983	**<0.001**
Action Type	3.520 (0.537 – 6.504)	1.522	2.313	**0.021**
log(Path)	0.882 (0.510 – 1.254)	0.190	4.650	**<0.001**
log(Step Number)	0.538 (-0.750 – 1.826)	0.657	0.819	0.413
Action Type[Tap] X log(Path)	-0.889 (-1.400 – -0.378)	0.261	-3.412	**0.001**
Action Type[Tap] X log(Step Number)	0.730 (-1.159 – 2.619)	0.964	0.757	0.449
log(Path) X log(Step Number)	-0.096 (-0.320 – 0.128)	0.114	-0.839	0.402
Action Type[Tap] X log(Path) X log(Step Number)	0.071 (-0.253 – 0.394)	0.165	0.430	0.668
Random Effects				
σ^2	0.34			
$\tau_{00\ \mathrm{subj}}$	4.09			
$\tau_{11\ \mathrm{subj.costChoiceTap}}$	6.70			
$\tau_{11\ \mathrm{subj.log(Path)}}$	0.11			
$\tau_{11\ \mathrm{subj.costChoiceTap:log(Path)}}$	0.19			
ρ_{01}	-0.88			
	-0.99			
	0.88			
ICC	0.16			
N_{subj}	40			
Observations	1440			
Marginal R^2 / Conditional R^2	0.245 / 0.367			
AIC	566.937			

Table A.3: Results of the mixed-effects linear regression model predicting movement durations by Action Type (Tapping vs. Dragging) * log(Path) * log(Step Number) in the Individual No Choice condition (the same model as reported in Table A.4.2., with a different reference level). Raw estimates on a log scale are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations. The reference level for Action Type is *Tapping*.

	Movement Duration ~ Action Type X log(Path) X log(Step Number) + (Action Type X log(Path) | participant)			
Predictors	*Estimates*	*std. Error*	*z*	*p*
Intercept	-1.935 (-3.982 – 0.113)	1.045	-1.852	0.064
Action Type	-3.520 (-6.504 – -0.537)	1.522	-2.313	**0.021**
log(Path)	-0.007 (-0.352 – 0.339)	0.176	-0.038	0.969
log(Step Number)	1.268 (-0.114 – 2.650)	0.705	1.798	0.072
Action Type[Drag] X log(Path)	0.889 (0.378 – 1.400)	0.261	3.412	**0.001**
Action Type[Drag] X log(Step Number)	-0.730 (-2.619 – 1.159)	0.964	-0.757	0.449
log(Path) X log(Step Number)	-0.025 (-0.258 – 0.208)	0.119	-0.209	0.834
Action Type[Drag] X log(Path) X log(Step Number)	-0.071 (-0.394 – 0.253)	0.165	-0.430	0.668
Random Effects				
σ^2	0.34			
$\tau_{00\ \text{subj}}$	4.09			
$\tau_{11\ \text{subj.costChoice1}}$	6.70			
$\tau_{11\ \text{subj.log(Path)}}$	0.11			
$\tau_{11\ \text{subj.costChoice1:log(Path)}}$	0.19			
ρ_{01}	-0.88			
	-0.99			
	0.88			
ICC	0.17			
N_{subj}	40			
Observations	1440			
Marginal R^2 / Conditional R^2	0.242 / 0.374			
AIC	566.937			

Table A.4: Results of the mixed-effects linear regression model predicting Movement Duration Indices by Tap Gain in the Individual No Choice condition. Raw estimates on a log scale are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Movement Duration Index ~ Tap Gain + (Tap Gain | participant)			
Predictors	*Estimates*	*std. Error*	*z*	*p*
Intercept	-1.109 (-1.233 – -0.985)	0.063	-17.546	**<0.001**
Tap Gain	0.013 (0.011 – 0.014)	0.001	18.671	**<0.001**
Random Effects				
σ^2	0.26			
$\tau_{00\ \text{subj}}$	0.13			
$\tau_{11\ \text{subj.TapGain}}$	0.00			
$\rho_{01\ \text{subj}}$	-0.86			
ICC	0.21			
N_{subj}	40			
Observations	1440			
Marginal R^2 / Conditional R^2	0.551 / 0.645			
AIC	1659.655			

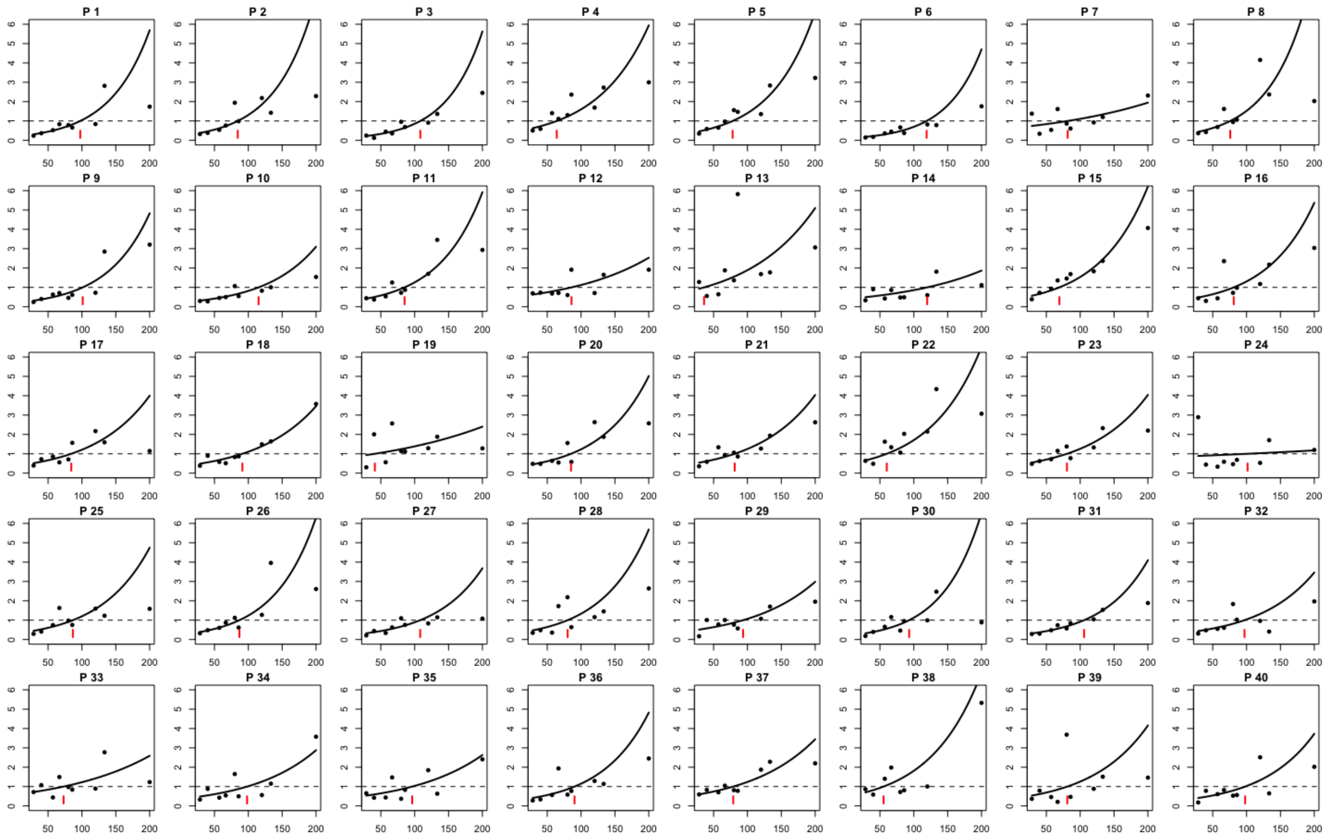

Figure A.1: Individual Movement Duration Indices (Movement Duration$_{Drag}$ / Movement Duration$_{Tap}$; y axis) in the Individual No Choice condition, plotted as a function of Tap Gain (Path / Step Number; x axis), with predicted regression lines based on the generalized linear mixed model Movement Duration Index ~ Tap Gain + (1 + Tap Gain | participant). The dashed horizontal line is at 1, signaling equivalence between the durations of dragging and tapping. Scatter points above this line indicate trials where tapping took a shorter time relative to dragging, points below the line come from trials where dragging took a shorter time relative to tapping. The red tick along the x axis indicates the threshold value of Tap Gain where each participant's dragging and tapping took an equal amount of time.

A.2 Individual Choice Condition

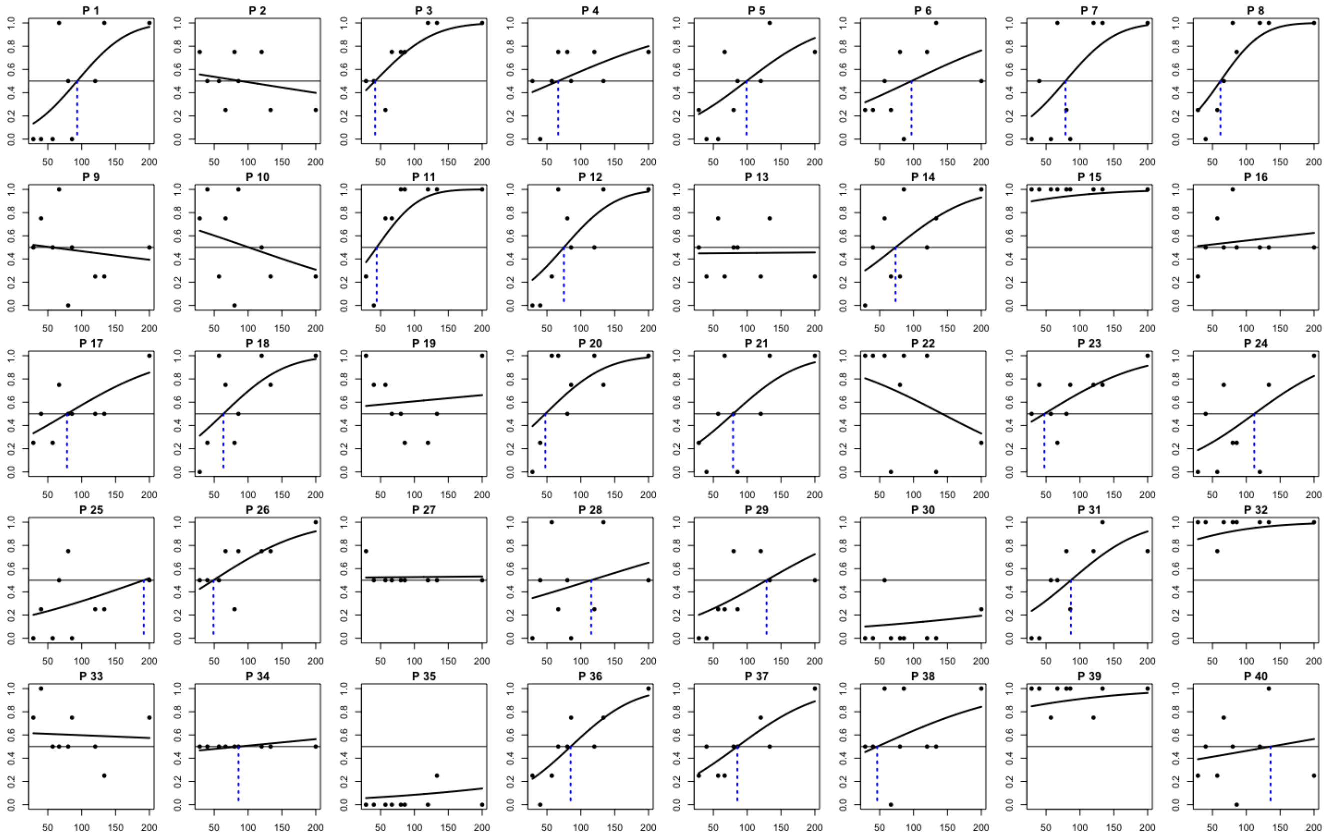

Figure A.2: Proportions of tap choices in the Individual Choice condition (y axis), plotted as a function of Tap Gain (Path / Step Number, x axis), with predicted regression lines based on the probit model p(Tap) ~ Tap Gain + (Tap Gain | participant). The horizontal line is at 50%, chance level for the 2AFC task. The vertical blue dashed lines along the x axis indicate the threshold value of Tap Gain where a participant was equally likely to choose dragging and tapping.

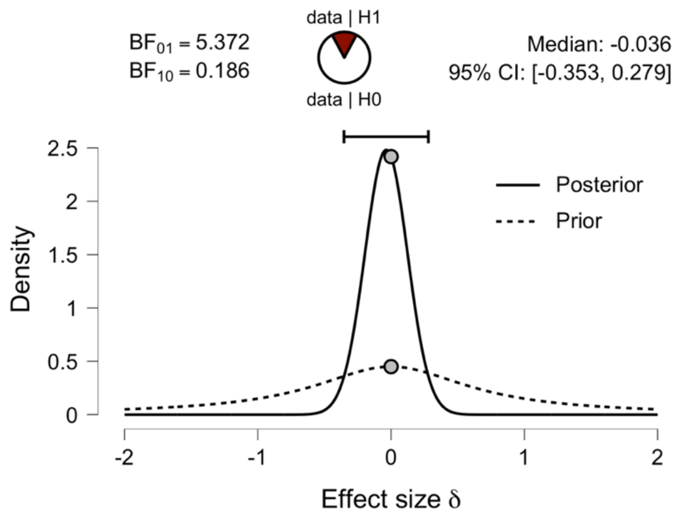

Figure A.3: Results of the Bayesian t-test comparing thresholds estimated in the Individual No Choice (Fig. A.4.1., a point of equal time benefit from dragging vs tapping based on Movement Duration Indices) and Choice conditions (Fig. A.4.2., PSE between tapping and dragging when choosing objects) against each other.

Table A.5: Results of the mixed effects probit regression model predicting the probability of Tapping in the Individual Choice condition by Tap Gain. Raw estimates on a log scale are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	p(Tap) ~ Tap Gain + (Tap Gain | participant)			
Predictors	*Log-Odds*	*std. Error*	*z*	*p*
Intercept	-0.496 (-0.800 – -0.191)	0.155	-3.192	**0.001**
Tap Gain	0.007 (0.004 – 0.010)	0.002	4.763	**<0.001**
Random Effects				
σ^2	1.00			
$\tau_{00\ \text{subj}}$	0.73			
$\tau_{11\ \text{subj.TapGain}}$	0.00			
$\rho_{01\ \text{subj}}$	-0.70			
ICC	0.36			
N_{subj}	40			
Observations	1440			
Marginal R^2 / Conditional R^2	0.081 / 0.416			
AIC	1753.261			

A.3 Joint Condition - Bayesian Logistic Regressions

Table A.6: Raw (pixel-based) parameter estimates and model fit measures (WAIC, LOO-CV, AUC – Area Under the Curve) of the logistic regression models fit to pooled data from *dyads with both PSE estimates, and participants with a PSE estimate, whose co-actors did not have a PSE (total n = 26).* Each row reports estimates for a model indicated in the first column. The best-fitting model is highlighted by bold font.

Model	μ_β Mode	μ_β 95% HDI	σ_β Mode	σ_β 95% HDI	WAIC [SE]	LOO-CV [SE]	AUC
1: Self Disparity	-0.004	-0.005, -0.002	0.004	0.003, 0.005	2528.3 [34.3]	2528.4 [34.3]	0.673
2: Other Disparity (egocentric)	-0.001	-0.002, -0.0003	0.001	0.001, 0.002	2871.2 [14.1]	2871.3 [14.1]	0.557
3: Self + Other Disparities (egocentric)	**Self: -0.004**	**-0.006, -0.003**	**0.004**	**0.003, 0.005**	**2400.3 [39.7]**	**2400.5 [39.7]**	**0.709**
	Other: -0.002	**-0.003, -0.001**	**0.002**	**0.001, 0.003**			

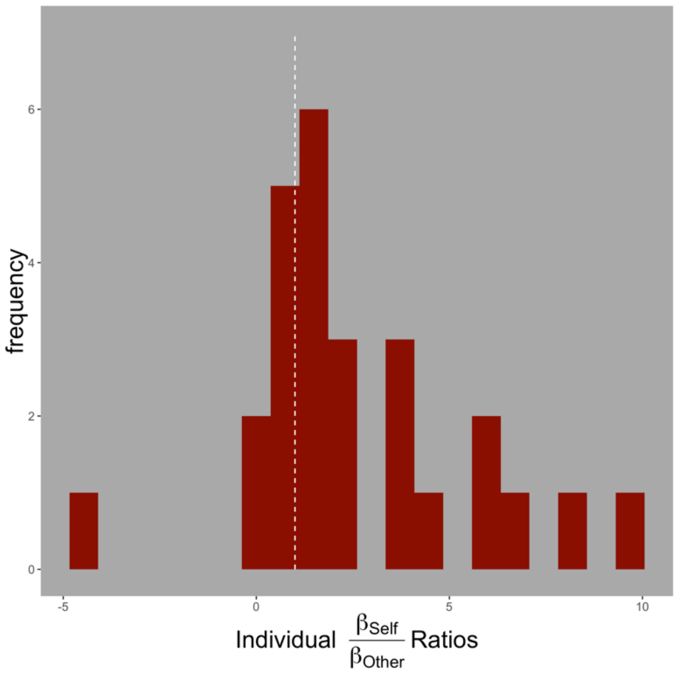

Figure A.4: Individual β_{Self} / β_{Other} ratios (n = 26) according to the best-fitting Self and egocentric Other Disparity model. The vertical white dashed line is at 1, indicating equal weights on Self and Other Disparities. Values below zero are results of the two β values having different signs: two participants had positive weights on Self, and negative weights on Other Disparities.

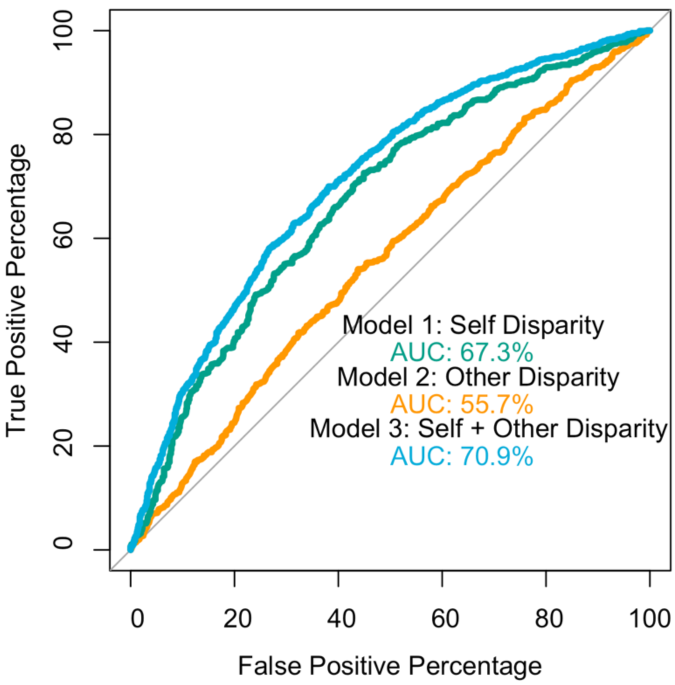

Figure A.5: Receiver Operating Characteristics (ROC) curves fitted to the three models tested on the pooled subsample's data (n = 26), with their respective Area Under the Curve values.

Table A.7: Raw (pixel-based) parameter estimates and model fit measures (WAIC, LOO-CV, AUC – Area Under the Curve) of the logistic regression models fit to data from *dyads with both PSE estimates* (n = 16). Each row reports estimates for a model indicated in the first column. The best-fitting model is highlighted by bold font.

Model	μ_β Mode	μ_β 95% HDI	σ_β Mode	σ_β 95% HDI	WAIC [SE]	LOO-CV [SE]	AUC
1: Self Disparity	-0.003	-0.005, -0.002	0.003	0.002, 0.005	1593.3 [26.2]	1593.3 [26.2]	0.681
2: Other Disparity (egocentric)	-0.001	-0.002, 3.08e-05	0.001	0.001, 0.003	1771.9 [10.6]	1771.9 [10.6]	0.548
3: Self + Other Disparity (egocentric)	**Self: -0.004**	**-0.006, -0.002**	**0.003**	**0.002, 0.005**	**1499.5 [31.5]**	**1499.6 [31.5]**	**0.723**
	Other: -0.002	**-0.004, -0.001**	**0.002**	**0.001, 0.003**			
4: Other Disparity (true)	-0.001	-0.002, -0.0003	0.001	0.001, 0.002	1766.3 [11.7]	1766.3 [11.7]	0.570
5: Self + Other Disparity (true)	Self: -0.004	-0.006, -0.002	0.003	0.002, 0.005	1541.2 [30.0]	1541.4 [30.0]	0.686
	Other: -0.001	-0.003, -0.0003	0.002	0.001, 0.003			

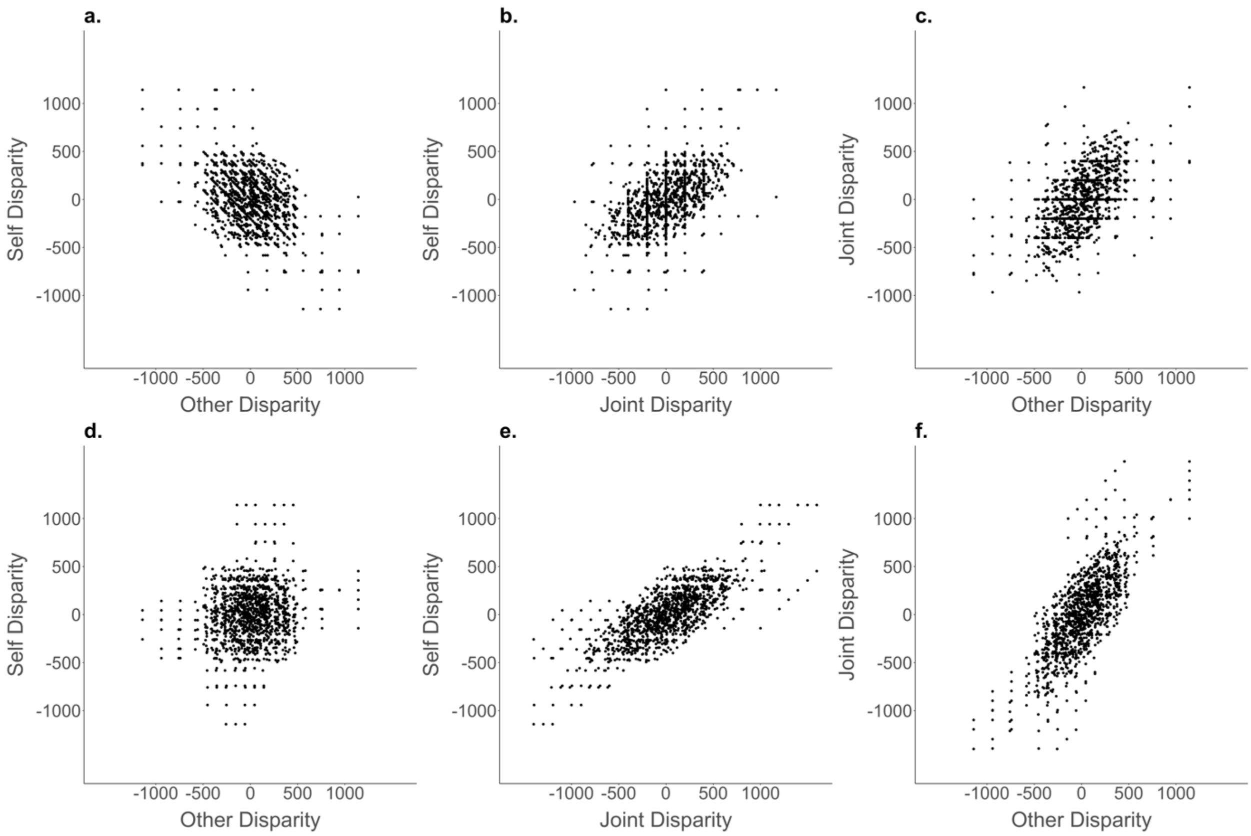

Figure A.6: Scatterplots of the joint distributions of cost disparities, collapsed across all trials of the complete dyad subsample (8 dyads, 1296 trials). (a) Self and egocentric Other Disparities were negatively correlated, $r = -.414$. (b) Self and egocentric Joint Disparities, and (c) egocentric Joint and Other Disparities were positively correlated with each other ($r = .541$ for both). (d) Self and true Other Disparities were positively correlated with each other, $r = .144$, as were (e) Self and true Joint Disparities, and (f) true Joint and Other Disparities ($r = .756$ for both).

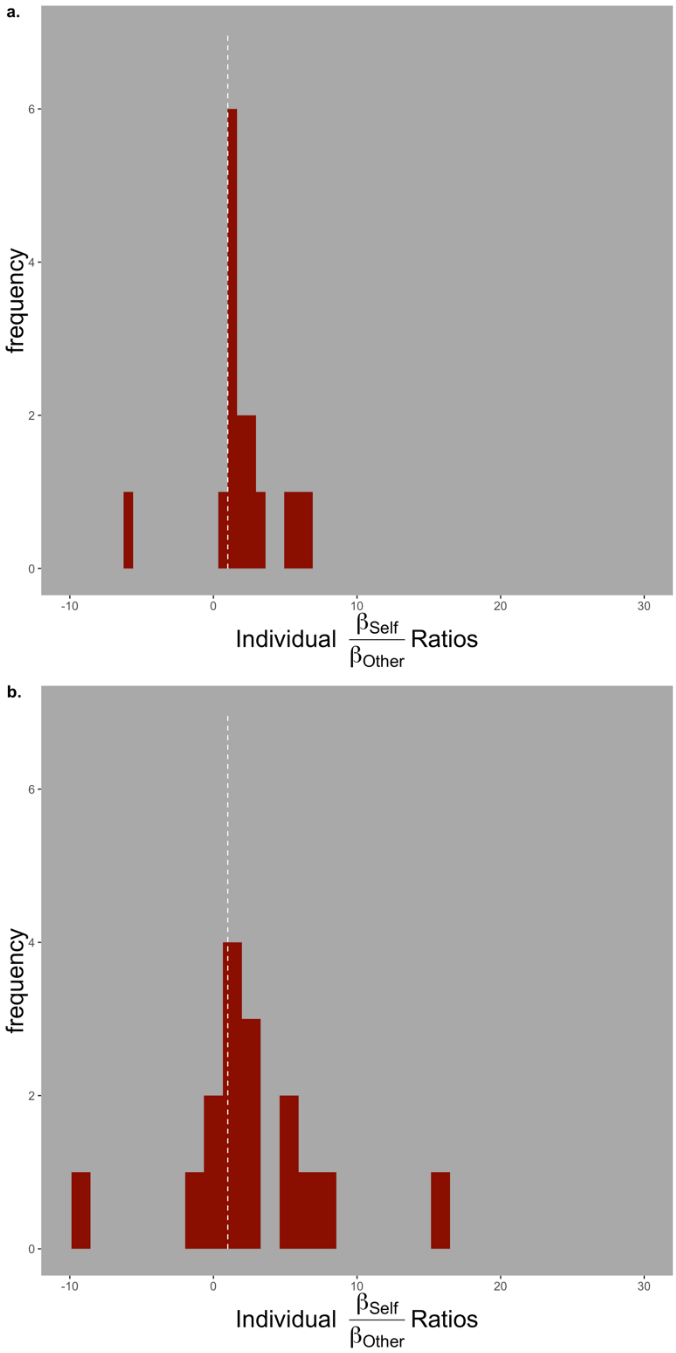

Figure A.7: Individual β_{Self} / β_{Other} ratios (n = 16) according to (a) the best-fitting Self and egocentric Other Disparity model and (b) the Self and true Other Disparity models. The vertical white dashed line is at 1, indicating equal weights on Self and Other Disparities. Values below zero are results of the two β values having different signs.

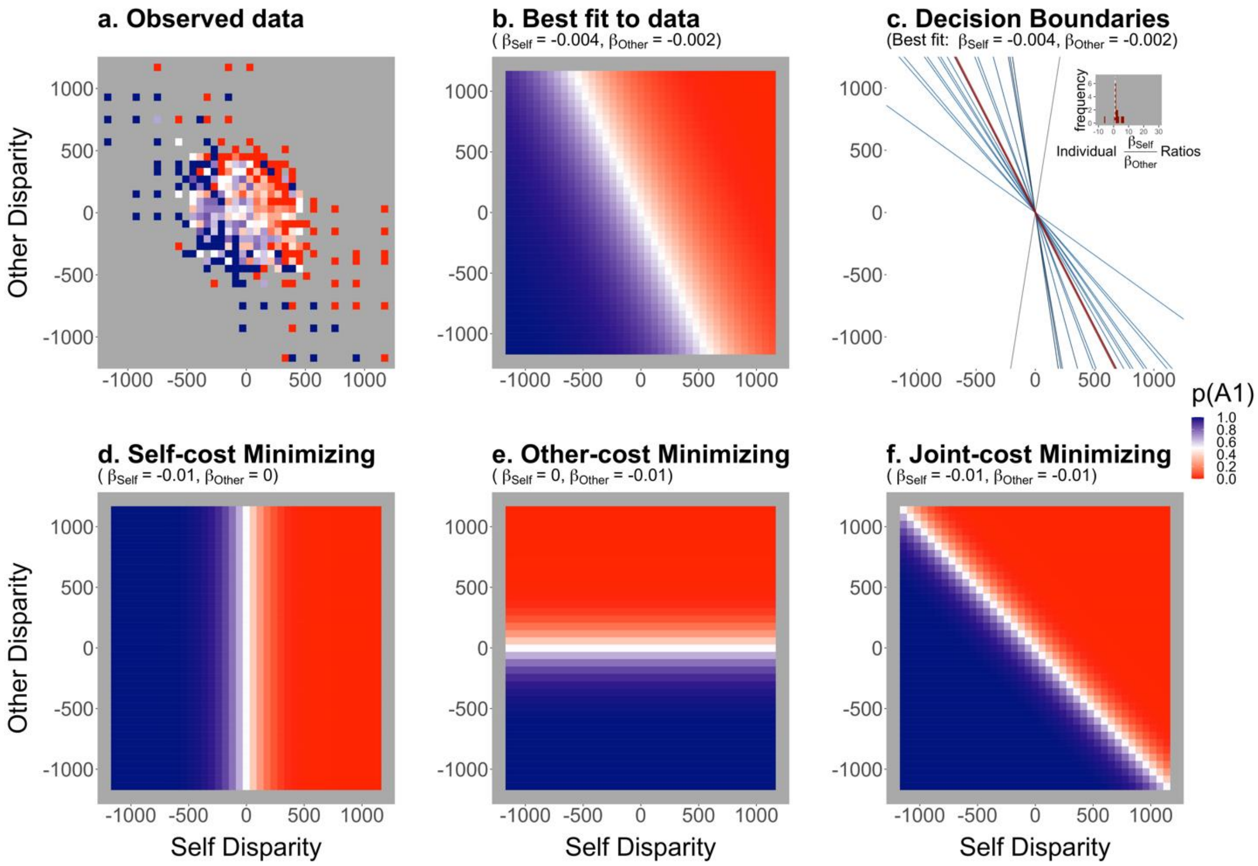

Figure A.8: (a) Observed object A1 choices (n = 16, bin width = 60), and (b) the posterior predictions of the best-fitting model using the linear combination of Self and egocentric Other Disparities. (c) Individual decision boundaries according to the best-fitting model; inset: frequency distribution of the $\mu_{\beta Self}/\mu_{\beta Other}$ ratios (see also Figure A.7a). The vertical white dashed line denotes 1, equal weights on Self and Other Disparity. (d-f) Predictions for optimal responses according to Self, egocentric Other, and Joint (i.e., Self + egocentric Other) cost-minimizing strategies, respectively. The lower the disparity to be minimized according to a model, the higher the probability of picking object A1 (blue). Predictions were calculated assuming that one pixel increase in a given parameter would result in 1% decrease in the odds of choosing object A1 over B1. All plots feature disparities in pixels.

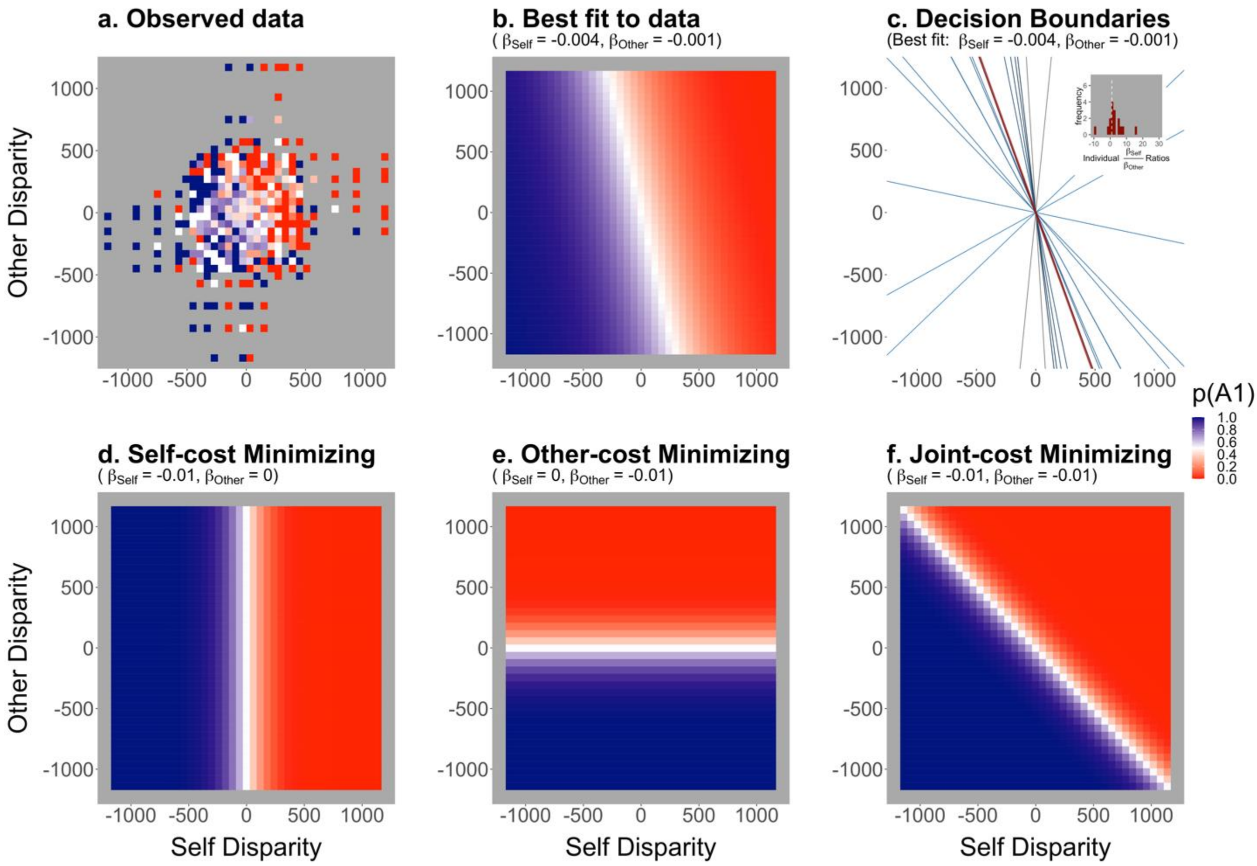

Figure A.9: (a) Observed object A1 choices (*n* = 16, bin width = 60), and (b) the posterior predictions of the model using the linear combination of Self and true Other Disparities. (c) Individual decision boundaries according to the model; inset: frequency distribution of the $\mu_{\beta Self}/\mu_{\beta Other}$ ratios (see also Figure A.7b). The vertical white dashed line denotes 1, equal weights on Self and Other Disparity. (d-f) Predictions for optimal responses according to Self, true Other, and Joint (i.e., Self + true Other) cost-minimizing strategies, respectively. The lower the disparity to be minimized according to a model, the higher the probability of picking object A1 (blue). Predictions were calculated assuming that one pixel increase in a given parameter would result in 1% decrease in the odds of choosing object A1 over B1. All plots feature disparities in pixels.

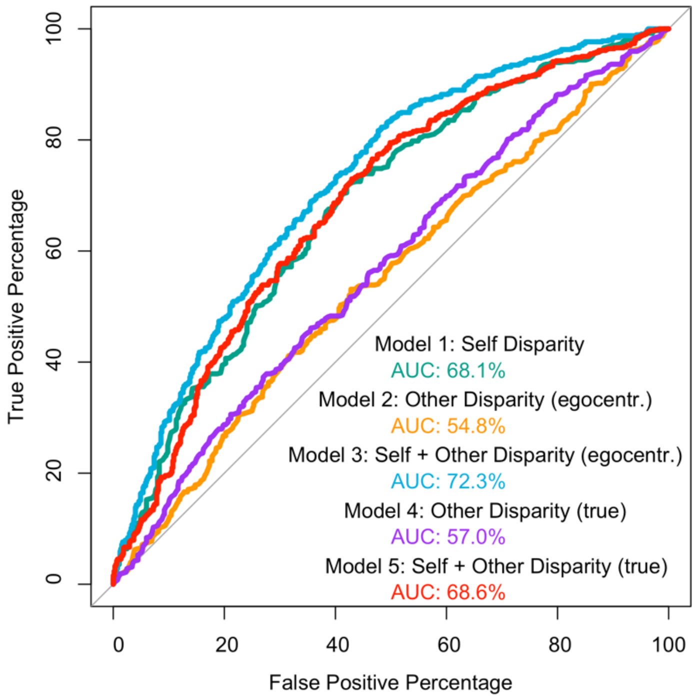

Figure A.10: Receiver Operating Characteristics (ROC) curves fitted to the five models tested on the complete dyads subsample's data (*n* = 16), with their respective Area Under the Curve values.

Table A.8: Raw (pixel-based) parameter estimates and model fit measures (WAIC, LOO-CV, AUC – Area Under the Curve) of the logistic regression models fit to data from *participants with a PSE estimate, whose co-actors did not have a PSE (n = 10)*. Each row reports estimates for a model indicated in the first column. The best-fitting model is highlighted by bold font.

Model	μ_β Mode	μ_β 95% HDI	σ_β Mode	σ_β 95% HDI	WAIC [SE]	LOO-CV [SE]	AUC
1: Self Disparity	-0.004	-0.008, -0.001	0.005	0.003, 0.009	935.3 [22.0]	935.3 [22.0]	0.656
2: Other Disparity (egocentric)	-0.001	-0.002, 0.0002	0.001	0.0004, 0.002	1101.1 [10.0]	1101.1 [10.0]	0.572
3: Self + Other Disparity (egocentric)	**Self: -0.005**	**-0.009, -0.001**	**0.005**	**0.003, 0.009**	**902.5 [24.7]**	**902.6 [24.7]**	**0.684**
	Other: -0.002	**-0.003, -0.0001**	**0.002**	**0.001, 0.003**			

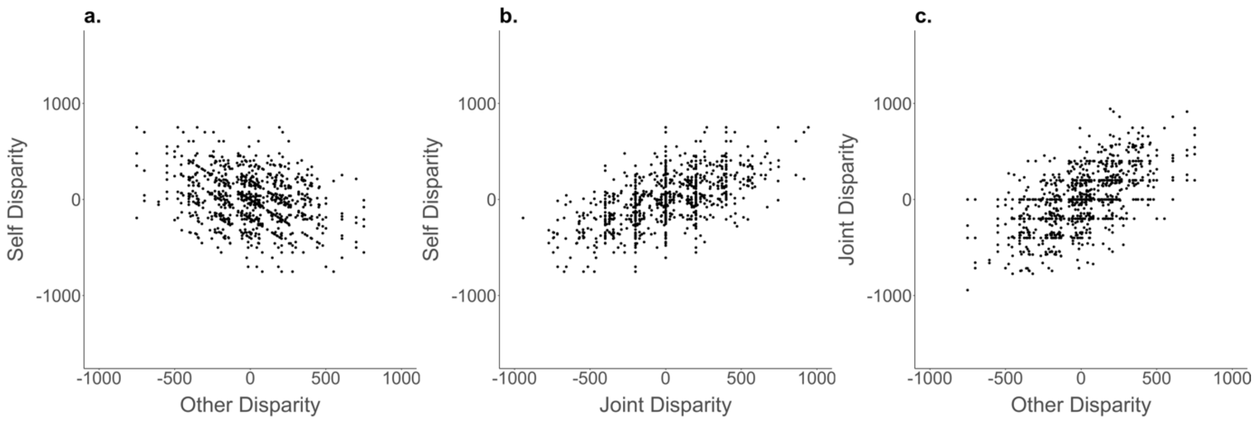

Figure A.11: Scatterplots of the joint distributions of cost disparities, collapsed across all trials of the partial dyad subsample (n = 10, 810 trials). (a) Self and Other Disparities were negatively correlated, r = -.414. (b) Self and Joint Disparities, and (c) Joint and Other Disparities were positively correlated with each other (both r = .541).

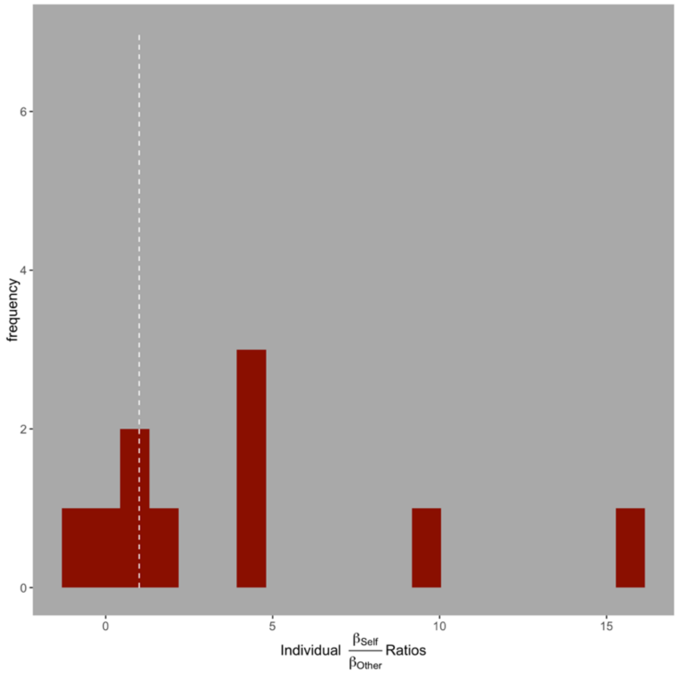

Figure A.12: Individual β_{Self} / β_{Other} ratios (n = 10) according to the best-fitting Self and egocentric Other Disparity model. The vertical white dashed line is at 1, indicating equal weights on Self and Other Disparities. Values below zero are results of the two β values having different signs.

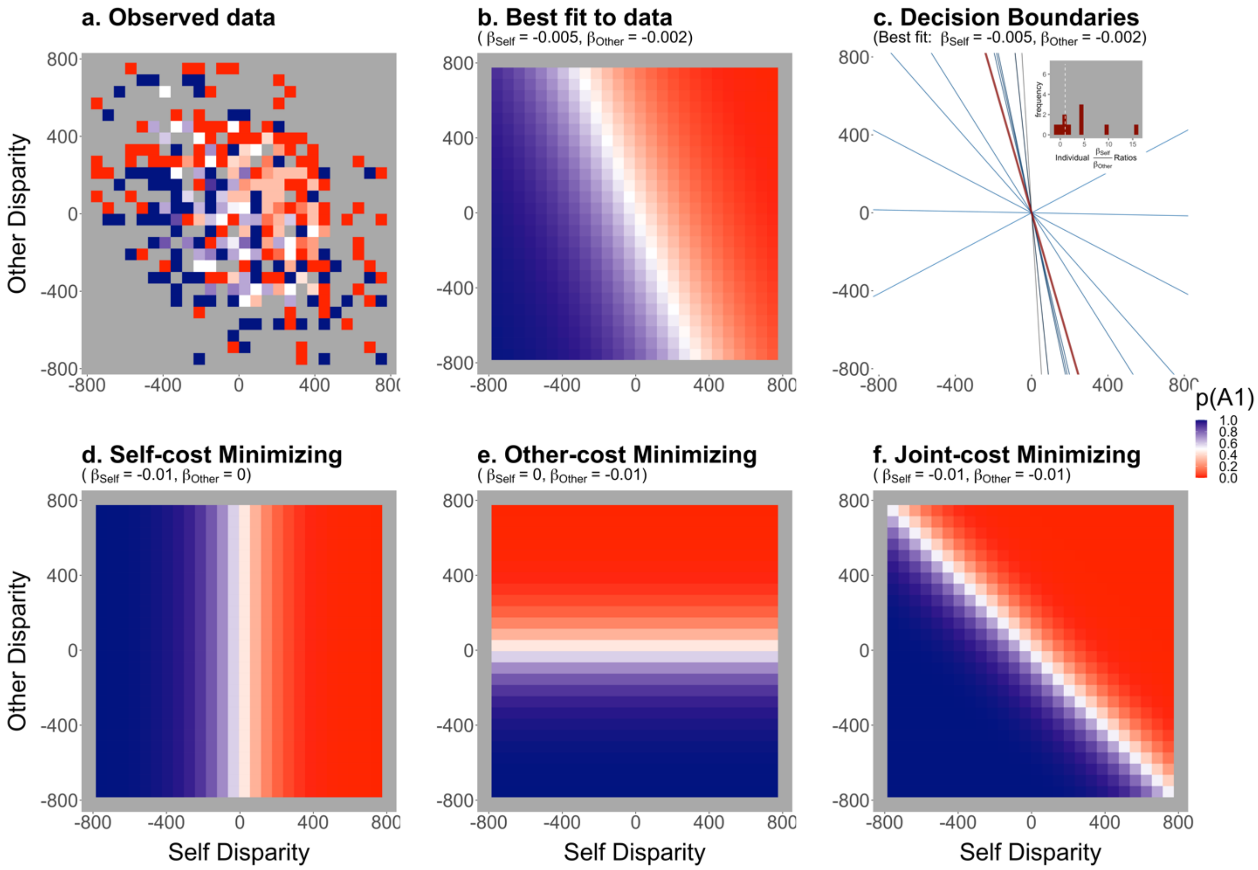

Figure A.13: (a) Observed object A1 choices (n = 10, bin width = 60), and (b) the posterior predictions of the best-fitting model using the linear combination of Self and egocentric Other Disparities. (c) Individual decision boundaries according to the best-fitting model; inset: frequency distribution of the $\mu_{\beta Self}/\mu_{\beta Other}$ ratios (see also Figure A.12). The vertical white dashed line denotes 1, equal weights on Self and Other Disparity. (d-f) Predictions for optimal responses according to Self, Other, and Joint (i.e., Self + Other) cost-minimizing strategies, respectively. The lower the disparity to be minimized according to a model, the higher the probability of picking object A1 (blue). Predictions were calculated assuming that one pixel increase in a given parameter would result in 1% decrease in the odds of choosing object A1 over B1. All plots feature disparities in pixels.

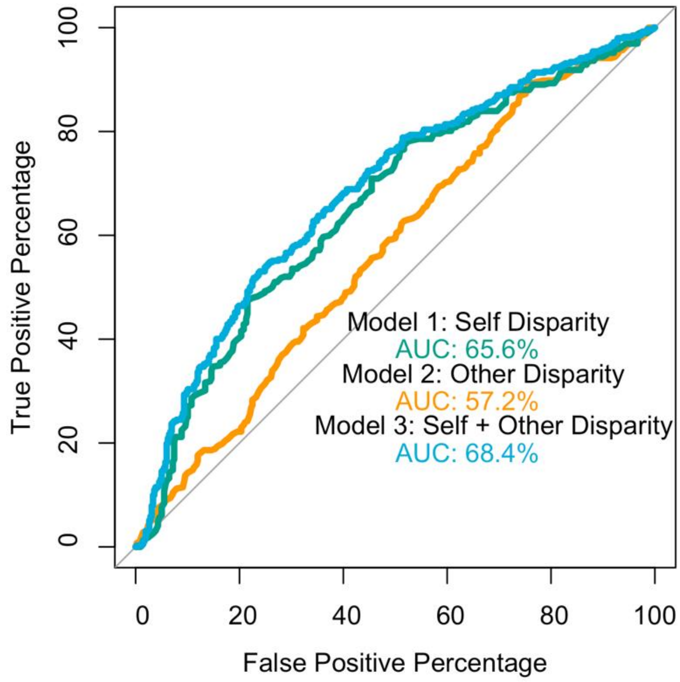

Figure A.14: Receiver Operating Characteristics (ROC) curves fitted to the three models tested on the partial dyads' data (n = 10), with their respective Area Under the Curve values.

A.4 Bayesian Model Structure

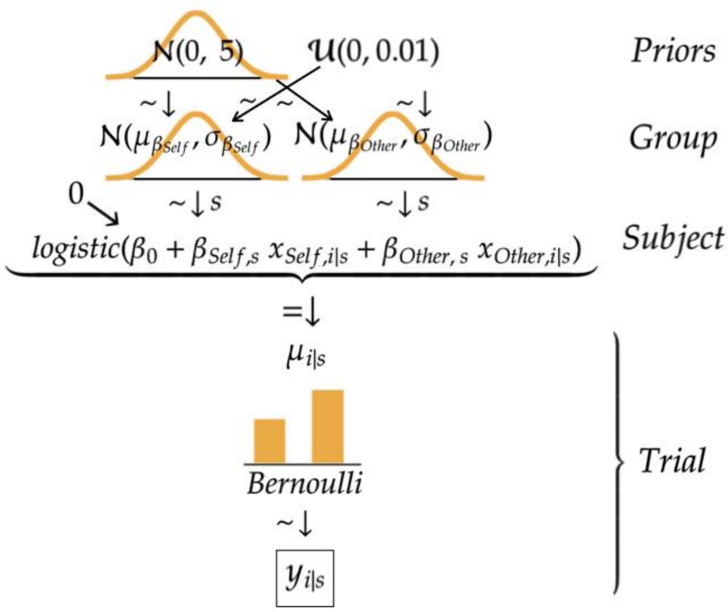

Figure A.15: A graphical schema of the hierarchical regression model, adapted from Kruschke (2015).

Appendix B – Additional Results to Chapter 5

B.1 Experiment 1

B.1.1 Additional information on congruency (section 5.2.4 Design)

Congruency is a factor we introduced in Török et al. (2019), reported in Chapter 2. In that study's task, the lengths of walls on the screen and the continuity of movement trajectories between co-actors ensured that in each trial, an individually efficient path choice for the participant was either also co-efficient (congruent trial) or not co-efficient (incongruent trial). In the present study, each trial could also be categorized either as congruent or incongruent in terms of the relationship between individually efficient and co-efficient action plans. If participants chose the object closer to them in the congruent trials, they chose an object that also minimized the total path length required from them and their partner as a group, were they to collect their successfully matched object choices (i.e., they were co-efficient as well as Self-cost minimizing). On the other hand, in the incongruent trials, collecting the object closer to the participant (Actor 1) would not minimize the total path length the dyad would have to move along to collect a matching object pair.

Due to the spatially separated movement trajectories of the co-actors, a subset of the congruent trials in the present study had an additional characteristic: namely, the individually efficient (Self-cost minimizing) object choice for the participant could sometimes be matched with an Other-cost minimizing movement on the partner's side. In these trials, therefore, Self-cost minimizing decisions would overlap with both Joint and Other-cost minimization ("joint-other congruent" trials). The remaining congruent trials would be classed as "joint congruent" if a Self-cost minimizing solution minimized also the joint costs of the trial but did *not* minimize the co-actor's individual costs (Other cost disparity). Note that these "joint congruent" trials are the same as the "congruent" trials in Chapter 2. The rest of the trials were incongruent, where choosing the Self-cost minimizing object was detrimental to co-efficiency.

Generally, overlaps in optimal solutions according to different cost-minimizing strategies make it harder to conclude about the cost-minimization strategy behind the decisions, therefore we aimed to minimize as much as possible the number of congruent trials generated in the present study.

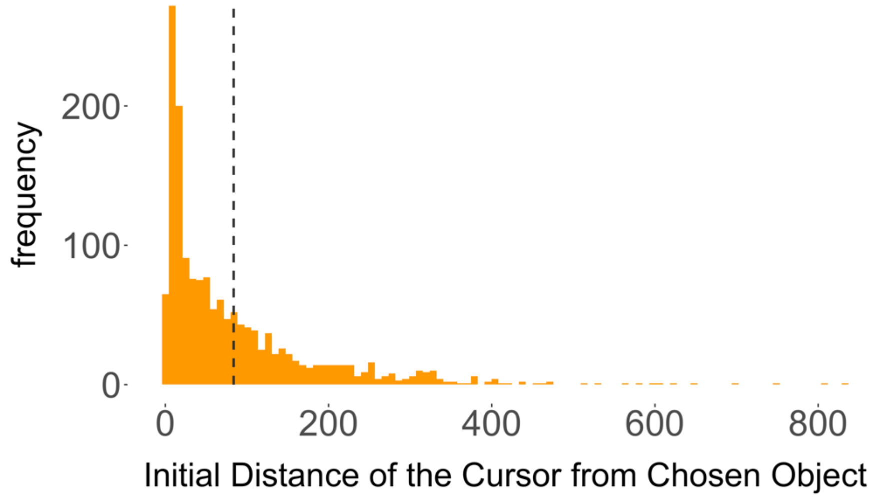

Figure B.1: Frequency histogram of the distances of the mouse cursor from the chosen object at the beginning of the Decision phase in Experiment 1. The dashed grey line indicates the mean of the distribution.

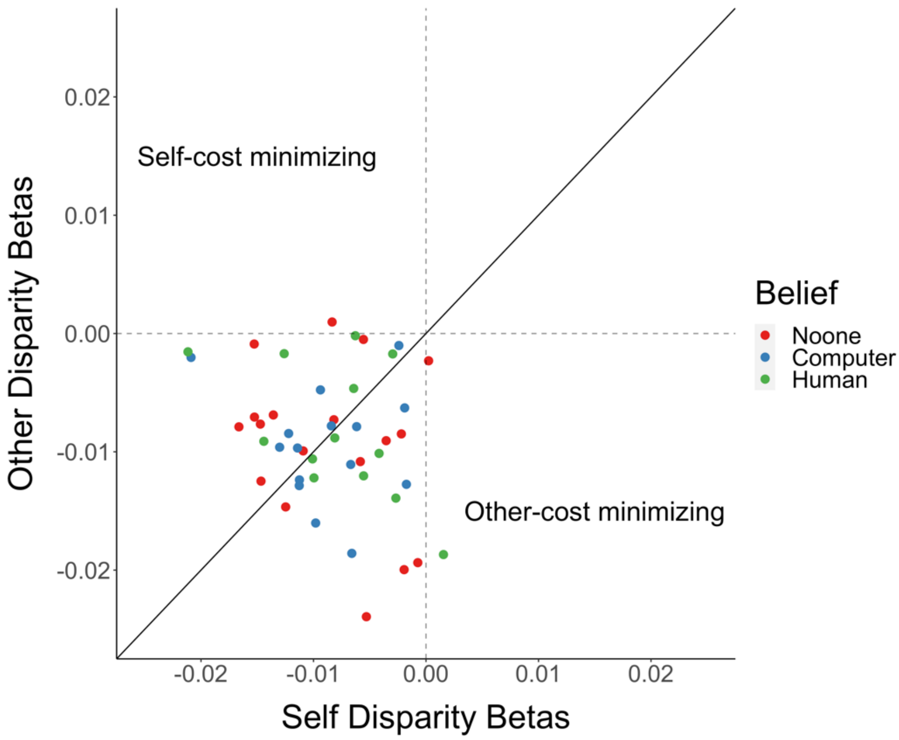

Figure B.2: Distribution of the estimated individual beta coefficients (N = 50) for the Self and Other Disparities in the best-fitting model (Model 5) in Experiment 1. The dashed grey lines indicate zero, above which disparities had a positive relationship with the probability of choosing the square, i.e., the larger the given disparity got, the higher the probability that the object was chosen. These values indicate that the given cost was not minimized by the participant's decision. The lower left region of the plot is where a joint-cost minimizing strategy would be displayed, and participants who placed similar weights on Self and Other Disparities are dispersed around the identity line. The dots are colored according to the participants' self-reported beliefs regarding the identity of the remote partner with whom they completed the task.

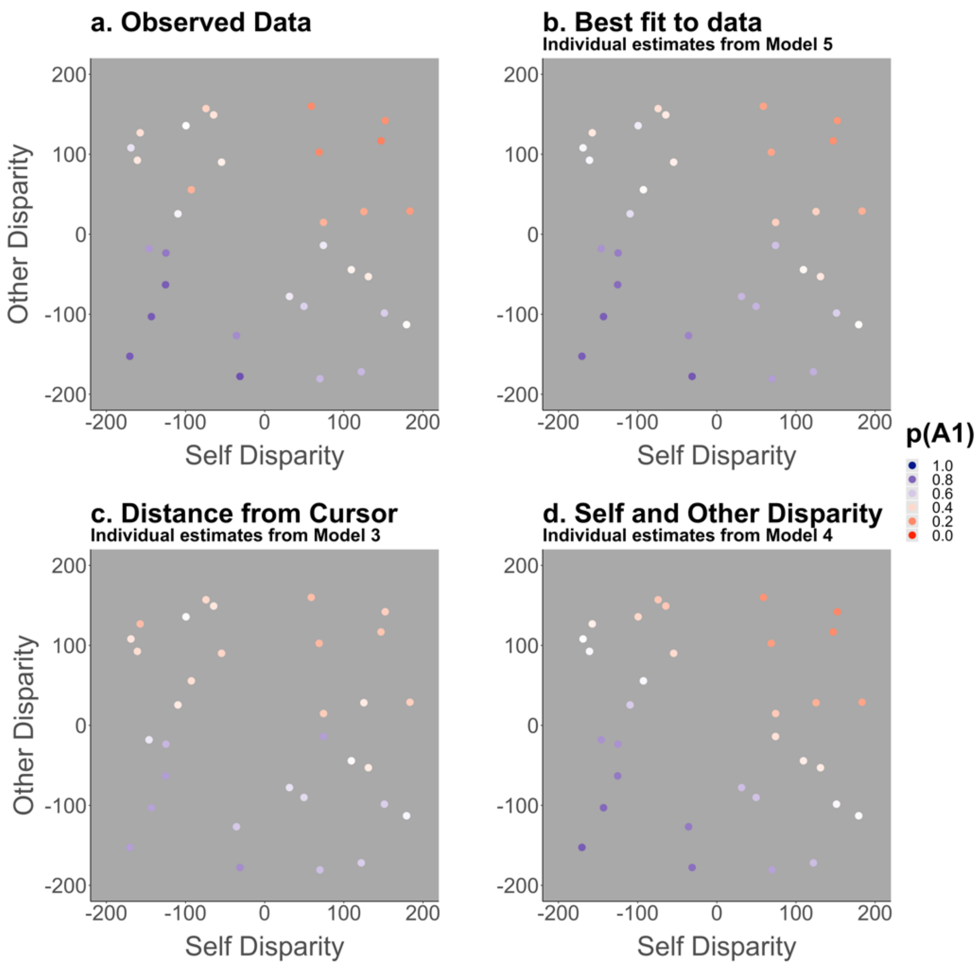

Figure B.3 (a) Observed object A1 choice probabilities for all trials in Experiment 1 ($N = 50$). Posterior predicted choice probabilities based on (b) the best-fitting model's mean of individual fixed effects estimates, including individual intercepts (Model 5, Self, Other Disparity and Distance from Cursor), the mean of individual estimates from the model including (c) only Distance from Cursor as predictor (Model 3), and (d) the linear combination of Self and Other Disparities (Model 4).

Table B.1: Results of Model 1, logistic regression predicting the probability of choosing the square object (A1) by Self Disparity only in Experiment 1. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 1			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	1.001 (0.656 – 1.528)	0.216	0.005	0.996
Self Disparity	0.995 (0.992 – 0.998)	0.002	-3.162	**0.002**
Random Effects				
σ^2	3.29			
$\tau_{00\ \text{subj}}$	2.08			
$\tau_{11\ \text{subj.Self\_Cost}}$	0.00			
$\rho_{01\ \text{subj}}$	-0.02			
ICC	0.51			
N_{subj}	50			
Observations	1585			
Marginal R^2 / Conditional R^2	0.047 / 0.537			
AIC	1803.848			

Table B.2: Results of Model 2, logistic regression predicting the probability of choosing the square object (A1) by Other Disparity only in Experiment 1. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 2			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	1.001 (0.656 – 1.527)	0.215	0.006	0.995
Other Disparity	0.992 (0.989 – 0.995)	0.002	-4.703	**<0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ \text{subj}}$	2.07			
$\tau_{11\ \text{subj.Other\_Cost}}$	0.00			
$\rho_{01\ \text{subj}}$	-0.00			
ICC	0.51			
N_{subj}	50			
Observations	1585			
Marginal R^2 / Conditional R^2	0.098 / 0.557			
AIC	1724.018			

Table B.3: Results of Model 3, logistic regression predicting the probability of choosing the square object (A1) by Distance from Cursor only in Experiment 1. Exponentiated estimates are reported with 95% confidence intervals are included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 3			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	50.913 (25.802 – 100.462)	0.347	11.333	**<0.001**
Distance from Cursor	0.972 (0.966 – 0.978)	0.003	-8.932	**<0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ subj}$	3.42			
$\tau_{11\ subj.dist\_A1}$	0.00			
$\rho_{01\ subj}$	-0.93			
ICC	0.72			
N_{subj}	50			
Observations	1585			
Marginal R^2 / Conditional R^2	0.515 / 0.864			
AIC	1050.585			

Table B.4: Results of Model 4, logistic regression predicting the probability of choosing the square object (A1) by Self and Other Disparities in Experiment 1. Exponentiated estimates are reported with 95% confidence intervals are included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 4			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	1.006 (0.604 – 1.675)	0.260	0.021	0.983
Self Disparity	0.992 (0.988 – 0.996)	0.002	-4.037	**<0.001**
Other Disparity	0.989 (0.985 – 0.993)	0.002	-5.151	**<0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ \mathrm{subj}}$	3.02			
$\tau_{11\ \mathrm{subj.Self\_Cost}}$	0.00			
$\tau_{11\ \mathrm{subj.Other\_Cost}}$	0.00			
ρ_{01}	0.00			
	0.01			
ICC	0.69			
N_{subj}	50			
Observations	1585			
Marginal R^2 / Conditional R^2	0.154 / 0.736			
AIC	1439.277			

Table B.5: Results of Model 5, logistic regression predicting the probability of choosing the square object (A1) by Self and Other Disparities and Distance from Cursor in Experiment 1. Exponentiated estimates are reported with 95% confidence intervals are included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 5			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	70.918 (30.884 – 162.844)	0.424	10.048	**<0.001**
Self Disparity	0.990 (0.986 – 0.994)	0.002	-5.092	**<0.001**
Other Disparity	0.992 (0.988 – 0.995)	0.002	-4.757	**<0.001**
Distance from Cursor	0.969 (0.961 – 0.977)	0.004	-7.411	**<0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ subj}$	0.56			
$\tau_{00\ subj.1}$	4.03			
$\tau_{11\ subj.Self\_Cost}$	0.00			
$\tau_{11\ subj.Other\_Cost}$	0.00			
$\tau_{11\ subj.1.dist\_A1}$	0.00			
$\rho_{01\ subj.Self\_Cost}$	0.32			
$\rho_{01\ subj.Other\_Cost}$	0.76			
$\rho_{01\ subj.1}$	-0.98			
ICC	0.50			
N_{subj}	50			
Observations	1585			
Marginal R^2 / Conditional R^2	0.741 / 0.870			
AIC	794.771			

Table B.6: Results of Model 6, logistic regression predicting the probability of choosing the square object (A1) by the categorical variable of whether the square was in the co-efficient or jointly sub-efficient position in a trial in Experiment 1. Exponentiated estimates are reported with 95% confidence intervals are included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 6			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	0.528 (0.332 – 0.838)	0.236	-2.708	**0.007**
Square is Co-efficient	3.483 (2.415 – 5.022)	0.187	6.681	**<0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ \mathrm{subj}}$	2.36			
$\tau_{11\ \mathrm{subj.SQcoeff1}}$	0.98			
$\rho_{01\ \mathrm{subj}}$	-0.45			
ICC	0.40			
N_{subj}	50			
Observations	1585			
Marginal R^2 / Conditional R^2	0.067 / 0.437			
AIC	1850.098			

Table B.7: Results of Model 7, logistic regression predicting the probability of choosing the square object (A1) by the Side of the screen on which the square was positioned in a trial in Experiment 1. Exponentiated estimates are reported with 95% confidence intervals are included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 7			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	1.148 (0.765 – 1.722)	0.207	0.665	0.506
Side[Right]	0.752 (0.536 – 1.054)	0.173	-1.655	0.098
Random Effects				
σ^2	3.29			
$\tau_{00\ \mathrm{subj}}$	1.79			
$\tau_{11\ \mathrm{subj.SQPos0.5}}$	0.81			
$\rho_{01\ \mathrm{subj}}$	-0.25			
ICC	0.37			
N_{subj}	50			
Observations	1585			
Marginal R^2 / Conditional R^2	0.004 / 0.369			
AIC	1977.662			

Table B.8: Results of Model 8, logistic regression predicting the probability of choosing the square object (A1) by Self Disparity and the Side of the screen where the square was positioned (-0.5 = Left, 0.5 = Right) in Experiment 1. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 8			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	1.355 (0.859 – 2.137)	0.232	1.306	0.191
Self Disparity	0.994 (0.990 – 0.997)	0.002	-3.743	**<0.001**
Side[Right]	0.558 (0.390 – 0.798)	0.183	-3.192	**0.001**
Self Disparity X Side[Right]	1.002 (0.999 – 1.004)	0.001	1.330	0.184
Random Effects				
σ^2	3.29			
$\tau_{00\ \mathrm{subj}}$	0.59			
$\tau_{00\ \mathrm{subj.1}}$	1.64			
$\tau_{11\ \mathrm{subj.Self\_Cost}}$	0.00			
$\tau_{11\ \mathrm{subj.1.SQPos0.5}}$	0.77			
$\rho_{01\ \mathrm{subj}}$	0.05			
$\rho_{01\ \mathrm{subj.1}}$	-0.13			
ICC	0.39			
N_{subj}	50			
Observations	1585			
Marginal R^2 / Conditional R^2	0.077 / 0.432			
AIC	1767.574			

Table B.9: Results of Model 9, logistic regression predicting the probability of choosing the square object (A1) by Other Disparity and the Side of the screen where the square was positioned (-0.5 = Left, 0.5 = Right) in Experiment 1. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 9			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	0.909 (0.542 – 1.525)	0.264	-0.360	0.719
Other Disparity	0.990 (0.986 – 0.994)	0.002	-5.516	**<0.001**
Side[Right]	1.074 (0.753 – 1.532)	0.181	0.394	0.694
Other Disparity X Side[Right]	1.005 (1.002 – 1.007)	0.001	3.424	**0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ subj}$	1.39			
$\tau_{00\ subj.1}$	1.51			
$\tau_{11\ subj.Other\_Cost}$	0.00			
$\tau_{11\ subj.1.SQPos0.5}$	0.70			
$\rho_{01\ subj}$	0.15			
$\rho_{01\ subj.1}$	-0.75			
ICC	0.46			
N_{subj}	50			
Observations	1585			
Marginal R^2 / Conditional R^2	0.123 / 0.524			
AIC	1709.769			

Table B.10: Results of Model 10, logistic regression predicting the probability of choosing the square object (A1) by Distance from Cursor and the Side of the screen where the square was positioned (-0.5 = Left, 0.5 = Right) in Experiment 1. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 10			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	54.384 (25.721 – 114.987)	0.382	10.460	**<0.001**
Distance from Cursor	0.971 (0.964 – 0.977)	0.003	-8.860	**<0.001**
Side[Right]	1.041 (0.533 – 2.033)	0.341	0.118	0.906
Distance from Cursor X Side[Right]	1.002 (0.998 – 1.005)	0.002	1.010	0.312
Random Effects				
σ^2	3.29			
$\tau_{00\ \text{subj}}$	3.21			
$\tau_{00\ \text{subj.1}}$	0.37			
$\tau_{11\ \text{subj.dist\_A1}}$	0.00			
$\tau_{11\ \text{subj.1.SQPos0.5}}$	0.39			
$\rho_{01\ \text{subj}}$	-0.99			
$\rho_{01\ \text{subj.1}}$	-0.08			
ICC	0.72			
N_{subj}	50			
Observations	1585			
Marginal R^2 / Conditional R^2	0.523 / 0.865			
AIC	1051.517			

Table B.11: Results of Model 11, logistic regression predicting the probability of choosing the square object (A1) by Self and Other Disparity and the Side of the screen where the square was positioned (-0.5 = Left, 0.5 = Right) in Experiment 1. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 11			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	1.042 (0.600 – 1.811)	0.282	0.146	0.884
Self Disparity	0.993 (0.989 – 0.997)	0.002	-3.349	**0.001**
Other Disparity	0.989 (0.985 – 0.993)	0.002	-5.042	**<0.001**
Side[Right]	0.885 (0.586 – 1.337)	0.211	-0.579	0.562
Self Disparity X Side[Right]	0.999 (0.996 – 1.002)	0.001	-0.761	0.446
Other Disparity X Side[Right]	1.001 (0.998 – 1.005)	0.002	0.775	0.439
Random Effects				
σ^2	3.29			
$\tau_{00\ subj}$	0.82			
$\tau_{00\ subj.1}$	2.39			
$\tau_{11\ subj.Self\_Cost}$	0.00			
$\tau_{11\ subj.Other\_Cost}$	0.00			
$\tau_{11\ subj.1.SQPos0.5}$	0.84			
$\rho_{01\ subj.Self\_Cost}$	0.09			
$\rho_{01\ subj.Other\_Cost}$	-0.07			
$\rho_{01\ subj.1}$	-0.14			
ICC	0.61			
N_{subj}	50			
Observations	1585			
Marginal R^2 / Conditional R^2	0.186 / 0.684			
AIC	1437.407			

Table B.12: Results of Model 12, logistic regression predicting the probability of choosing the square object (A1) by Self and Other Disparity, the Distance from Cursor and the Side of the screen where the square was positioned (-0.5 = Left, 0.5 = Right) in Experiment 1. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 12			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	53.720 (20.019 – 144.152)	0.504	7.910	**<0.001**
Self Disparity	0.991 (0.986 – 0.996)	0.003	-3.468	**0.001**
Other Disparity	0.990 (0.986 – 0.994)	0.002	-4.524	**<0.001**
Distance from Cursor	0.970 (0.962 – 0.978)	0.004	-7.166	**<0.001**
Side[Right]	1.574 (0.555 – 4.462)	0.532	0.853	0.393
Self Disparity X Side[Right]	0.997 (0.992 – 1.002)	0.002	-1.231	0.218
Other Disparity X Side[Right]	1.001 (0.996 – 1.007)	0.003	0.435	0.663
Dist. from Cursor X Side[Right]	1.000 (0.995 – 1.005)	0.002	0.065	0.948
Random Effects				
σ^2	3.29			
τ_{00} subj	1.78			
τ_{00} subj.1	1.55			
τ_{00} subj.2	1.77			
τ_{11} subj.Self_Cost	0.00			
τ_{11} subj.Other_Cost	0.00			
τ_{11} subj.1.dist_A1	0.00			
τ_{11} subj.2.SQPos0.5	1.63			
ρ_{01} subj.Self_Cost	0.54			
ρ_{01} subj.Other_Cost	0.62			
ρ_{01} subj.1	-1.00			
ρ_{01} subj.2	0.17			
ICC	0.60			
N subj	50			
Observations	1585			
Marginal R^2 / Conditional R^2	0.689 / 0.877			
AIC	815.372			

Table B.13: Results of Model 13, logistic regression predicting the probability of choosing the square object (A1) by a linear combination of Self Disparity and Distance from Cursor in Experiment 1. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 13			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	62.768 (30.760 – 128.083)	0.364	11.375	**<0.001**
Self Disparity	0.993 (0.989 – 0.997)	0.002	-3.765	**<0.001**
Distance from Cursor	0.971 (0.964 – 0.977)	0.003	-8.670	**<0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ \mathrm{subj}}$	3.24			
$\tau_{11\ \mathrm{subj.Self\_Cost}}$	0.00			
$\tau_{11\ \mathrm{subj.dist\_A1}}$	0.00			
ρ_{01}	0.24			
	-0.90			
ICC	0.78			
N_{subj}	50			
Observations	1585			
Marginal R^2 / Conditional R^2	0.490 / 0.888			
AIC	898.745			

Table B.14: Results of Model 14, logistic regression predicting the probability of choosing the square object (A1) by a linear combination of Other Disparity and Distance from Cursor in Experiment 1. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 14			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	46.688 (22.213 – 98.130)	0.379	10.141	**<0.001**
Other Disparity	0.995 (0.992 – 0.998)	0.001	-3.257	**0.001**
Distance from Cursor	0.972 (0.966 – 0.979)	0.004	-8.045	**<0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ \text{subj}}$	4.20			
$\tau_{11\ \text{subj.Other\_Cost}}$	0.00			
$\tau_{11\ \text{subj.dist\_A1}}$	0.00			
ρ_{01}	0.26			
	-0.92			
ICC	0.77			
N_{subj}	50			
Observations	1585			
Marginal R^2 / Conditional R^2	0.483 / 0.880			
AIC	995.733			

Table B.15: Results of Model 15, logistic regression predicting the probability of choosing the square object (A1) by the intercept only in Experiment 1. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 15			
Predictors	Odds Ratios	std. Error	z	p
Intercept	0.995 (0.689 – 1.439)	0.188	-0.025	0.980
Random Effects				
σ^2	3.29			
$\tau_{00\ subj}$	1.58			
ICC	0.32			
N_{subj}	50			
Observations	1585			
Marginal R^2 / Conditional R^2	0.000 / 0.324			
AIC	2003.071			

Table B.16: Summary of all the logistic regression models predicting the probability of choosing the square object (A1) in Experiment 1. No. of Par$_i$ = number of estimated parameters for model *i*; log(L$_i$) = natural logarithm of the maximum likelihood for model *i*; AIC$_i$ = Akaike Information Criterion for model *i*; w_i(AIC) = the rounded Akaike weights; Marg. R^2 = marginal R-squared of the fixed effects; Cond. R^2 = conditional R-squared of the fixed and random effects combined. Both were calculated following Nakagawa et al.'s method (2017), using the *performance* package in R. "dist_A1" means Distance of object A1 from Cursor, "SQCoeff" signifies the categorical variable of Square being the object in the co-efficient position in a trial (1) or in the jointly sub-efficient position (0). "Side" is the factor of whether the square object is on the left (-0.5) or right (0.5) side of the screen for the participant. Grouping factor N = 50 (subj), total number of observations = 1585 for all models.

Model ID	**Formula: Object Choice (1 = Square) ~**	**No. of Par$_i$**	**log(L$_i$)**	**AIC$_i$**	**w_i(AIC)**	**Marg. R^2**	**Cond. R^2**
M1	~ Self_Cost + (1 + Self_Cost | subj)	5	-896.9	1803.8	7.6^{e-220}	.047	.537
M2	~ Other_Cost + (1 + Other_Cost | subj)	5	-857	1724	1.6^{e-202}	.098	. 557
M3	~ dist_A1 + (1 + dist_A1 | subj)	5	-520.3	1050.6	2.8^{e-56}	.515	. 864
M4	~ Self_Cost + Other_Cost + (1 + Self_Cost + Other_Cost | subj)	9	-710.6	1439.3	1.1^{e-140}	.154	.736
M5	~ Self_Cost + Other_Cost + dist_A1 + (1 + Self_Cost + Other_Cost | subj) + (1 + dist_A1 | subj)	13	-384.4	794.8	**0.999**	**.741**	. 870
M6	~SQcoeff + (1 + SQcoeff| subj)	5	-920	1850.1	6.9^{e-230}	.067	.437
M7	~ Side + (1 + Side | subj)	5	-983.8	1977.7	1.4^{e-257}	.004	.369
M8	~ Self_Cost * Side + (1 + Self_Cost | subj) + (1 + Side | subj)	10	-873.8	1767.6	5.7^{e-212}	.077	.432
M9	~ Other_Cost * Side + (1 + Other_Cost | subj) + (1 + Side | subj)	10	-844.9	1709.8	2^{e-199}	.123	.524
M10	~ dist_A1 * Side + (1 + dist_A1 | subj) + (1 + Side | subj)	10	-515.8	1051.5	1.8^{e-56}	.523	.865
M11	~ (Self_Cost + Other_Cost)*Side + (1 + Self_Cost + Other_Cost | subj) + (1 + Side | subj)	15	-703.7	1437.4	2.8^{e-140}	.186	.684
M12	~ (Self_Cost + Other_Cost + dist_A1) * Side + (1	20	-387.7	815.4	**3.4^{e-05}**	.689	.877

	+ Self_Cost + Other_Cost |subj) + (1 + dist_A1 | subj) + (1 + Side | subj)						
M13	~ Self_Cost + dist_A1 + (1 + Self_Cost + dist_A1 | subj)	9	-440.4	898.7	2.6^{e-23}	.490	.888
M14	~ Other_Cost + dist_A1 + (1 + Other_Cost + dist_A1 | subj)	9	-488.9	995.7	2.3^{e-44}	.483	.880
M15	Intercept only: ~ (1 | subj)	2	-999.5	2003.1	4.2^{e-263}	.000	.324

B.2 Experiment 2

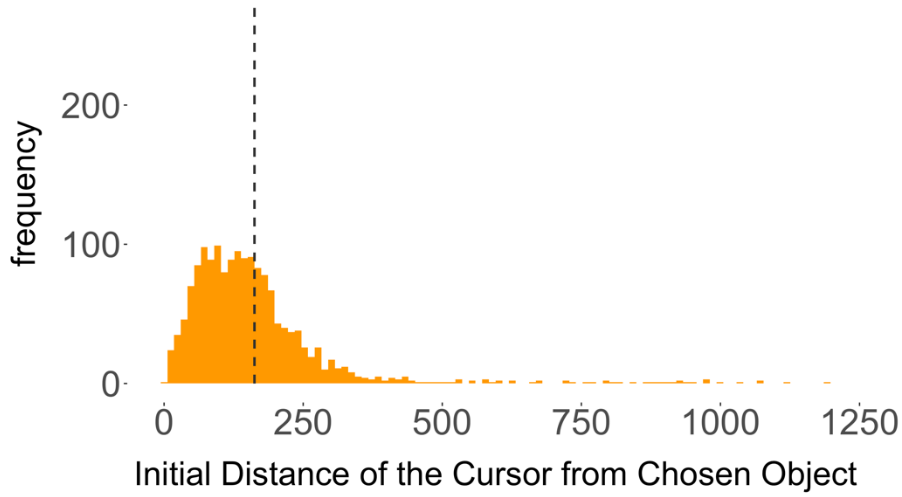

Figure B.4: Frequency histogram of the distances of the mouse cursor from the chosen object at the beginning of the Decision phase in Experiment 2. The dashed grey line indicates the mean of the distribution.

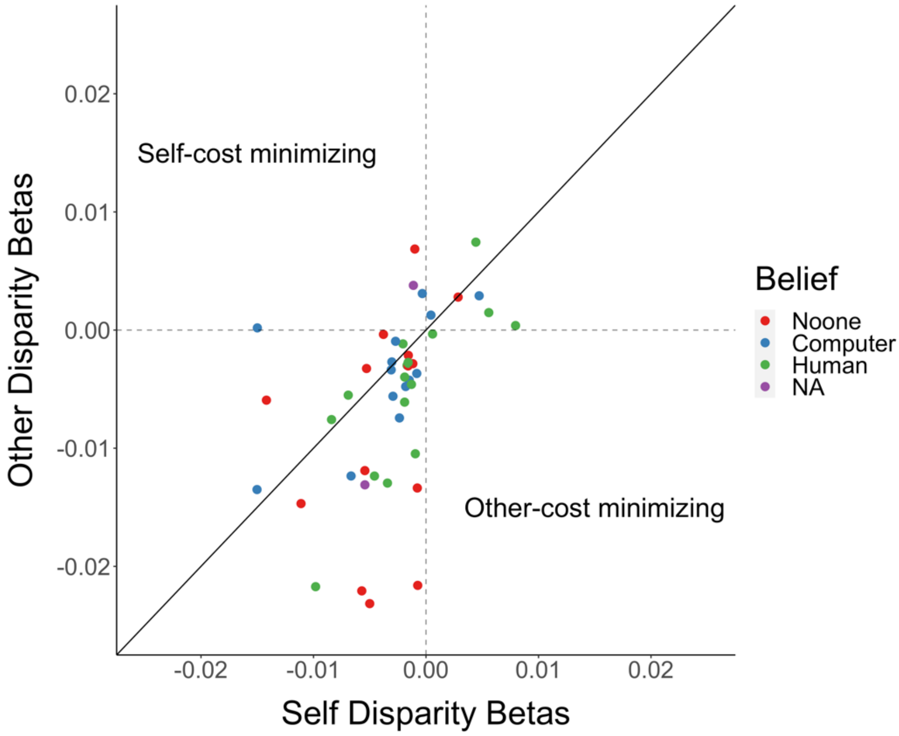

Figure B.5: Distribution of the estimated individual beta coefficients (N = 50) for the Self and Other Disparities in the best-fitting model (Model 5) in Experiment 2. The dashed grey lines indicate zero, above which disparities had a positive relationship with the probability of choosing the square, i.e., the larger the given cost disparity got, the higher the probability that the object was chosen. These values indicate that the given cost was not minimized by the participant's decision. The lower left region of the plot is where a joint-cost minimizing strategy would be displayed, and participants who placed similar weights on Self and Other Disparities are dispersed around the identity line. The dots are colored according to the participants' self-reported beliefs regarding the identity of the remote partner with whom they completed the task (with 2 NAs).

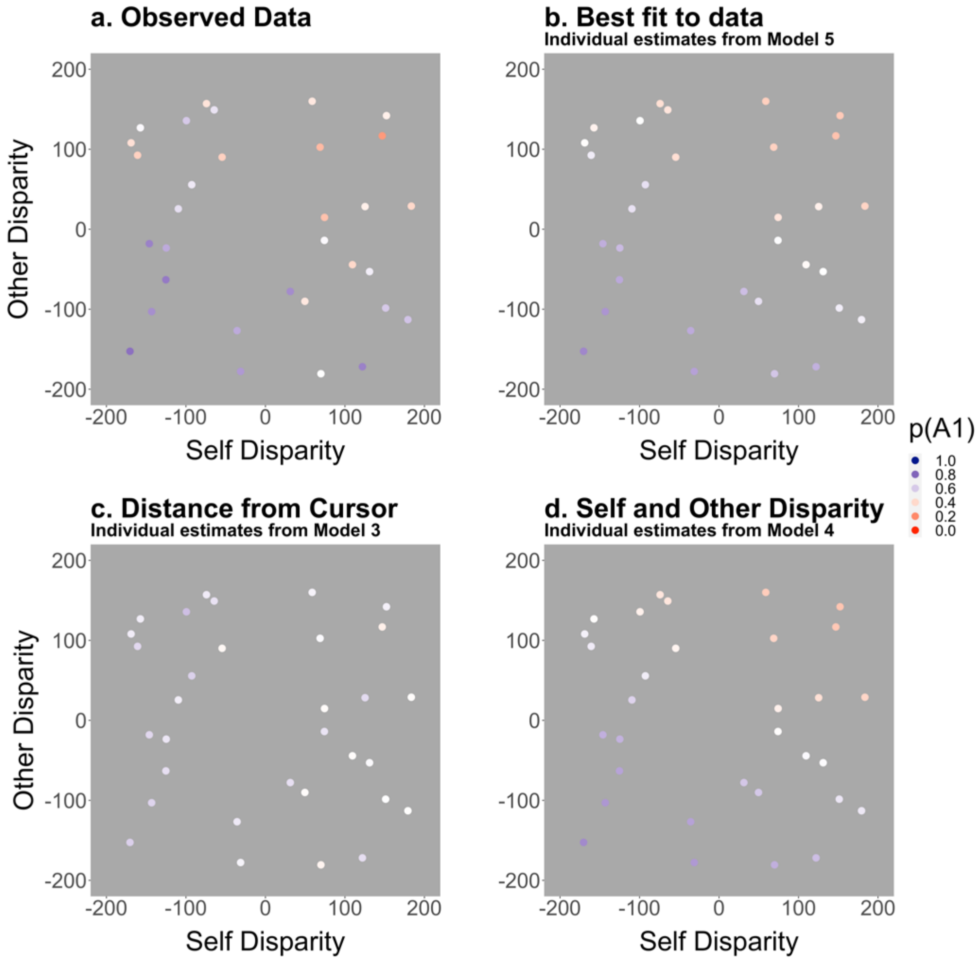

Figure B.6: (a) Observed object A1 choice probabilities for all trials in Experiment 2 ($N = 50$). Posterior predicted choice probabilities based on (b) the best-fitting model's mean of individual fixed effects estimates, including individual intercepts (predictors: Self, Other Disparity and Distance from Cursor), and based on the mean of individual estimates from the model including (c) only Distance from Cursor as predictor (Model 3), and (d) the linear combination of Self and Other Disparities (Model 4).

Table B.17: Results of Model 1, logistic regression predicting the probability of choosing the square object (A1) by Self Disparity only in Experiment 2. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 1			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	1.235 (0.796 – 1.917)	0.224	0.942	0.346
Self Disparity	0.998 (0.996 – 1.000)	0.001	-2.512	**0.012**
Random Effects				
σ^2	3.29			
$\tau_{00\ \mathrm{subj}}$	2.28			
$\tau_{11\ \mathrm{subj.Self\_Cost}}$	0.00			
$\rho_{01\ \mathrm{subj}}$	-0.04			
ICC	0.44			
N_{subj}	50			
Observations	1590			
Marginal R^2 / Conditional R^2	0.012 / 0.451			
AIC	1912.538			

Table B.18: Results of Model 2, logistic regression predicting the probability of choosing the square object (A1) by Other Disparity only in Experiment 2. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 2			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	1.243 (0.784 – 1.972)	0.235	0.926	0.354
Other Disparity	0.995 (0.993 – 0.998)	0.001	-3.393	**0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ \mathrm{subj}}$	2.51			
$\tau_{11\ \mathrm{subj.Other\_Cost}}$	0.00			
$\rho_{01\ \mathrm{subj}}$	0.09			
ICC	0.51			
N_{subj}	50			
Observations	1590			
Marginal R^2 / Conditional R^2	0.037 / 0.524			
AIC	1822.640			

Table B.19: Results of Model 3, logistic regression predicting the probability of choosing the square object (A1) by Distance from Cursor only in Experiment 2. Exponentiated estimates are reported with 95% confidence intervals are included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 3			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	3.018 (1.901 – 4.791)	0.236	4.682	**<0.001**
Distance from Cursor	0.994 (0.992 – 0.997)	0.001	-4.728	**<0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ subj}$	1.82			
$\tau_{11\ subj.dist\_A1}$	0.00			
$\rho_{01\ subj}$	-0.28			
ICC	0.49			
N_{subj}	50			
Observations	1590			
Marginal R^2 / Conditional R^2	0.087 / 0.536			
AIC	1890.919			

Table B.20: Results of Model 4, logistic regression predicting the probability of choosing the square object (A1) by Self and Other Disparities in Experiment 2. Exponentiated estimates are reported with 95% confidence intervals are included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 4			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	1.259 (0.763 – 2.077)	0.255	0.900	0.368
Self Disparity	0.996 (0.994 – 0.999)	0.001	-3.191	**0.001**
Other Disparity	0.994 (0.991 – 0.997)	0.002	-3.603	**<0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ subj}$	2.97			
$\tau_{11\ subj.Self\_Cost}$	0.00			
$\tau_{11\ subj.Other\_Cost}$	0.00			
ρ_{01}	-0.00			
	0.08			
ICC	0.59			
N_{subj}	50			
Observations	1590			
Marginal R^2 / Conditional R^2	0.059 / 0.612			
AIC	1729.134			

Table B.21: Results of Model 5, logistic regression predicting the probability of choosing the square object (A1) by Self and Other Disparities and Distance from Cursor in Experiment 2. Exponentiated estimates are reported with 95% confidence intervals are included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 5			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	3.002 (1.763 – 5.112)	0.272	4.049	**<0.001**
Self Disparity	0.997 (0.995 – 0.999)	0.001	-2.729	**0.006**
Other Disparity	0.994 (0.990 – 0.997)	0.002	-3.808	**<0.001**
Distance from Cursor	0.995 (0.992 – 0.997)	0.001	-4.108	**<0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ \mathrm{subj}}$	1.53			
$\tau_{00\ \mathrm{subj.1}}$	0.84			
$\tau_{11\ \mathrm{subj.Self\_Cost}}$	0.00			
$\tau_{11\ \mathrm{subj.Other\_Cost}}$	0.00			
$\tau_{11\ \mathrm{subj.1.dist\_A1}}$	0.00			
$\rho_{01\ \mathrm{subj.Self\_Cost}}$	-0.01			
$\rho_{01\ \mathrm{subj.Other\_Cost}}$	-0.02			
$\rho_{01\ \mathrm{subj.1}}$	-0.23			
ICC	0.49			
N_{subj}	50			
Observations	1590			
Marginal R^2 / Conditional R^2	0.148 / 0.569			
AIC	1685.475			

Table B.22: Results of Model 6, logistic regression predicting the probability of choosing the square object (A1) by the categorical variable of whether the square was in the co-efficient or jointly sub-efficient position in a trial in Experiment 2. Exponentiated estimates are reported with 95% confidence intervals are included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 6			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	0.832 (0.510 – 1.358)	0.250	-0.734	0.463
Square is Co-efficient	2.198 (1.498 – 3.225)	0.196	4.025	**<0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ \mathrm{subj}}$	2.70			
$\tau_{11\ \mathrm{subj.SQcoeff1}}$	1.14			
$\rho_{01\ \mathrm{subj}}$	-0.37			
ICC	0.44			
N_{subj}	50			
Observations	1590			
Marginal R^2 / Conditional R^2	0.026 / 0.457			
AIC	1885.504			

Table B.23: Results of Model 7, logistic regression predicting the probability of choosing the square object (A1) by the Side of the screen on which the square was positioned in a trial in Experiment 2. Exponentiated estimates are reported with 95% confidence intervals are included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 7			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	1.461 (0.960 – 2.225)	0.214	1.769	0.077
Side[Right]	0.705 (0.541 – 0.917)	0.135	-2.599	**0.009**
Random Effects				
σ^2	3.29			
$\tau_{00\ \mathrm{subj}}$	1.93			
$\tau_{11\ \mathrm{subj.SQPos0.5}}$	0.26			
$\rho_{01\ \mathrm{subj}}$	0.14			
ICC	0.40			
N_{subj}	50			
Observations	1590			
Marginal R^2 / Conditional R^2	0.006 / 0.400			
AIC	1959.155			

Table B.24: Results of Model 8, logistic regression predicting the probability of choosing the square object (A1) by Self Disparity and the Side of the screen where the square was positioned (-0.5 = Left, 0.5 = Right) in Experiment 2. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 8			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	1.639 (1.046 – 2.569)	0.229	2.156	**0.031**
Self Disparity	0.996 (0.994 – 0.998)	0.001	-3.552	**<0.001**
Side[Right]	0.595 (0.454 – 0.781)	0.138	-3.749	**<0.001**
Self Disparity X Side[Right]	1.002 (1.000 – 1.004)	0.001	1.953	0.051
Random Effects				
σ^2	3.29			
$\tau_{00\ \mathrm{subj}}$	1.12			
$\tau_{00\ \mathrm{subj.1}}$	1.08			
$\tau_{11\ \mathrm{subj.Self\_Cost}}$	0.00			
$\tau_{11\ \mathrm{subj.1.SQPos0.5}}$	0.22			
$\rho_{01\ \mathrm{subj}}$	0.06			
$\rho_{01\ \mathrm{subj.1}}$	0.31			
ICC	0.31			
N_{subj}	50			
Observations	1590			
Marginal R^2 / Conditional R^2	0.030 / 0.332			
AIC	1895.012			

Table B.25: Results of Model 9, logistic regression predicting the probability of choosing the square object (A1) by Other Disparity and the Side of the screen where the square was positioned (-0.5 = Left, 0.5 = Right) in Experiment 2. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 9			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	1.346 (0.826 – 2.194)	0.249	1.194	0.232
Other Disparity	0.996 (0.993 – 0.998)	0.001	-3.026	**0.002**
Side[Right]	0.863 (0.645 – 1.154)	0.148	-0.996	0.319
Other Disparity X Side[Right]	1.000 (0.997 – 1.002)	0.001	-0.082	0.935
Random Effects				
σ^2	3.29			
$\tau_{00\ \mathrm{subj}}$	0.97			
$\tau_{00\ \mathrm{subj.1}}$	1.66			
$\tau_{11\ \mathrm{subj.Other\_Cost}}$	0.00			
$\tau_{11\ \mathrm{subj.1.SQPos0.5}}$	0.28			
$\rho_{01\ \mathrm{subj}}$	0.12			
$\rho_{01\ \mathrm{subj.1}}$	-0.17			
ICC	0.36			
N_{subj}	50			
Observations	1590			
Marginal R^2 / Conditional R^2	0.049 / 0.389			
AIC	1828.259			

Table B.26: Results of Model 10, logistic regression predicting the probability of choosing the square object (A1) by Distance from Cursor and the Side of the screen where the square was positioned (-0.5 = Left, 0.5 = Right) in Experiment 2. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 10			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	3.761 (2.231 – 6.342)	0.267	4.970	**<0.001**
Distance from Cursor	0.994 (0.992 – 0.997)	0.001	-4.286	**<0.001**
Side[Right]	0.682 (0.457 – 1.020)	0.205	-1.865	0.062
Distance from Cursor X Side[Right]	1.000 (0.998 – 1.001)	0.001	-0.422	0.673
Random Effects				
σ^2	3.29			
$\tau_{00\ subj}$	1.04			
$\tau_{00\ subj.1}$	0.93			
$\tau_{11\ subj.dist\_A1}$	0.00			
$\tau_{11\ subj.1.SQPos0.5}$	0.10			
$\rho_{01\ subj}$	-0.23			
$\rho_{01\ subj.1}$	-1.00			
ICC	0.47			
N_{subj}	50			
Observations	1590			
Marginal R^2 / Conditional R^2	0.098 / 0.518			
AIC	1885.727			

Table B.27: Results of Model 11, logistic regression predicting the probability of choosing the square object (A1) by Self and Other Disparity and the Side of the screen where the square was positioned (-0.5 = Left, 0.5 = Right) in Experiment 2. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 11			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	1.528 (0.892 – 2.617)	0.275	1.542	0.123
Self Disparity	0.996 (0.993 – 0.998)	0.001	-3.286	**0.001**
Other Disparity	0.995 (0.992 – 0.999)	0.002	-2.760	**0.006**
Side[Right]	0.751 (0.551 – 1.024)	0.158	-1.811	0.070
Self Disparity X Side[Right]	1.001 (0.998 – 1.003)	0.001	0.526	0.599
Other Disparity X Side[Right]	0.998 (0.995 – 1.001)	0.001	-1.552	0.121
Random Effects				
σ^2	3.29			
$\tau_{00\ subj}$	0.67			
$\tau_{00\ subj.1}$	1.22			
$\tau_{00\ subj.2}$	1.29			
$\tau_{11\ subj.Self\_Cost}$	0.00			
$\tau_{11\ subj.1.Other\_Cost}$	0.00			
$\tau_{11\ subj.2.SQPos0.5}$	0.29			
$\rho_{01\ subj}$	-0.15			
$\rho_{01\ subj.1}$	0.14			
$\rho_{01\ subj.2}$	-0.22			
ICC	0.28			
N_{subj}	50			
Observations	1590			
Marginal R^2 / Conditional R^2	0.103 / 0.351			
AIC	1734.696			

Table B.28: Results of Model 12, logistic regression predicting the probability of choosing the square object (A1) by Self and Other Disparity, the Distance from Cursor and the Side of the screen where the square was positioned (-0.5 = Left, 0.5 = Right) in Experiment 2. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 12			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	3.132 (1.752 – 5.598)	0.296	3.854	**<0.001**
Self Disparity	0.996 (0.994 – 0.999)	0.001	-2.856	**0.004**
Other Disparity	0.995 (0.991 – 0.998)	0.002	-2.988	**0.003**
Distance from Cursor	0.995 (0.993 – 0.998)	0.001	-3.239	**0.001**
Side[Right]	0.949 (0.570 – 1.580)	0.260	-0.200	0.841
Self Disparity X Side[Right]	1.001 (0.998 – 1.003)	0.001	0.531	0.595
Other Disparity X Side[Right]	0.998 (0.995 – 1.001)	0.001	-1.361	0.174
Dist. from Cursor X Side[Right]	0.999 (0.997 – 1.001)	0.001	-1.031	0.303
Random Effects				
σ^2	3.29			
$\tau_{00\ subj}$	0.55			
$\tau_{00\ subj.1}$	1.04			
$\tau_{00\ subj.2}$	0.64			
$\tau_{11\ subj.Self\_Cost}$	0.00			
$\tau_{11\ subj.Other\_Cost}$	0.00			
$\tau_{11\ subj.1.dist\_A1}$	0.00			
$\tau_{11\ subj.2.SQPos0.5}$	0.32			
$\rho_{01\ subj.Self\_Cost}$	-0.04			
$\rho_{01\ subj.Other\_Cost}$	-0.06			
$\rho_{01\ subj.1}$	-0.11			
$\rho_{01\ subj.2}$	-0.27			
ICC	0.41			
N_{subj}	50			
Observations	1590			
Marginal R^2 / Conditional R^2	0.166 / 0.505			
AIC	1692.002			

Table B.29: Results of Model 13, logistic regression predicting the probability of choosing the square object (A1) by a linear combination of Self Disparity and Distance from Cursor in Experiment 2. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 13			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	2.507 (1.822 – 3.450)	0.163	5.647	**<0.001**
Self Disparity	0.998 (0.997 – 1.000)	0.001	-1.779	0.075
Distance from Cursor	0.996 (0.993 – 0.998)	0.001	-3.479	**0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ \mathrm{subj}}$	0.43			
$\tau_{11\ \mathrm{subj.Self\_Cost}}$	0.00			
$\tau_{11\ \mathrm{subj.dist\_A1}}$	0.00			
ρ_{01}	-0.05			
	0.17			
ICC	0.53			
N_{subj}	50			
Observations	1590			
Marginal R^2 / Conditional R^2	0.059 / 0.554			
AIC	1871.330			

Table B.30: Results of Model 14, logistic regression predicting the probability of choosing the square object (A1) by a linear combination of Other Disparity and Distance from Cursor in Experiment 2. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 14			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	3.536 (2.094 – 5.972)	0.267	4.724	**<0.001**
Other Disparity	0.995 (0.992 – 0.998)	0.001	-3.674	**<0.001**
Distance from Cursor	0.993 (0.991 – 0.996)	0.001	-4.651	**<0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ \mathrm{subj}}$	2.34			
$\tau_{11\ \mathrm{subj.Other\_Cost}}$	0.00			
$\tau_{11\ \mathrm{subj.dist\_A1}}$	0.00			
ρ_{01}	-0.13			
	-0.32			
ICC	0.61			
N_{subj}	50			
Observations	1590			
Marginal R^2 / Conditional R^2	0.119 / 0.655			
AIC	1739.595			

Table B.31: Results of Model 15, logistic regression predicting the probability of choosing the square object (A1) by the intercept only in Experiment 2. Exponentiated estimates are reported with 95% confidence intervals included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

	Model 15			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	1.228 (0.811 – 1.858)	0.211	0.971	0.332
Random Effects				
σ^2	3.29			
$\tau_{00\ \mathrm{subj}}$	2.02			
ICC	0.38			
N_{subj}	50			
Observations	1590			
Marginal R^2 / Conditional R^2	0.000 / 0.381			
AIC	1967.994			

Table B.32: Summary of all the logistic regression models predicting the probability of choosing the square object (A1) in Experiment 2. No. of Par_i = number of estimated parameters for model *i*; $\log(L_i)$ = natural logarithm of the maximum likelihood for model *i*; AIC_i = Akaike Information Criterion for model *i*; $w_i(AIC)$ = the rounded Akaike weights; Marg. R^2 = marginal R-squared of the fixed effects; Cond. R^2 = conditional R-squared of the fixed and random effects combined. Both were calculated following Nakagawa et al.'s method (2017), using the *performance* package in R. "dist_A1" means Distance of object A1 from Cursor, "SQCoeff" signifies the categorical variable of Square being the object in the co-efficient position in a trial (1) or in the jointly sub-efficient position (0). "Side" is the factor of whether the square object is on the left (-0.5) or right (0.5) side of the screen for the participant. Grouping factor N = 50 (subj), total number of observations = 1590 for all models.

Model ID	**Formula: Object Choice (1 = Square) ~**	**No. of Par_i**	**$\log(L_i)$**	**AIC_i**	**$w_i(AIC)$**	**Marg. R^2**	**Cond. R^2**
M1	~ Self_Cost + (1 + Self_Cost | subj)	5	-951.3	1912.5	4.8^{e-50}	.012	.451
M2	~ Other_Cost + (1 + Other_Cost | subj)	5	-906.3	1822.6	1.6^{e-30}	.037	.524
M3	~ dist_A1 + (1 + dist_A1 | subj)	5	-940.5	1890.9	2.4^{e-45}	.087	.536
M4	~ ObjChoice ~ Self_Cost + Other_Cost + (1 + Self_Cost + Other_Cost | subj)	9	-855.6	1729.1	3.2^{e-10}	.059	.612
M5	~ Self_Cost + Other_Cost + dist_A1 + (1 + Self_Cost + Other_Cost | subj) + (1 + dist_A1 | subj)	13	-829.7	1685.5	**0.963**	.148	.569
M6	~SQcoeff + (1 + SQcoeff| subj)	5	-937.8	1885.5	3.5^{e-44}	.026	.457
M7	~ Side + (1 + Side | subj)	5	-974.6	1959.2	3.6^{e-60}	.006	.400
M8	~ Self_Cost * Side + (1 + Self_Cost | subj) + (1 + Side | subj)	10	-937.5	1895	3^{e-46}	.030	.332
M9	~ Other_Cost * Side + (1 + Other_Cost | subj) + (1 + Side | subj)	10	-904.1	1828.3	9.5^{e-32}	.049	.389
M10	~ dist_A1 * Side + (1 + dist_A1 | subj) + (1 + Side | subj)	10	-932.9	1885.7	3.2^{e-44}	.098	.518
M11	~ (Self_Cost + Other_Cost)*Side + (1 + Self_Cost | subj) + (1 + Other_Cost | subj) + (1 + Side | subj)	15	-852.3	1734.7	2^{e-11}	.103	.351
M12	~ (Self_Cost + Other_Cost + dist_A1) * Side + (1 + Self_Cost + Other_Cost |subj) + (1 + dist_A1 | subj) + (1 + Side | subj)	20	-826	1692	**0.037**	**.166**	.505
M13	~ Self_Cost + dist_A1 + (1 + Self_Cost + dist_A1 | subj)	9	-926.7	1871.3	4.2^{e-41}	.059	.554
M14	~ Other_Cost + dist_A1 + (1 + Other_Cost + dist_A1 | subj)	9	-860.8	1739.6	1.7^{e-12}	.119	.655
M15	Intercept only: ~ (1 | subj)	2	-982	1968	4.3^{e-62}	.000	.381

B.3 Comparison between Experiments – Logistic Regression models

Table B.33: Summary of the logistic regression models run to compare the joint-cost minimization effect between experiments.

Model ID	Formula: Object Choice (1 = Square) ~	No. of Par_i	$\log(L_i)$	AIC_i	$w_i(AIC)$	Marg. R^2	Cond. R^2
M5	~ Self_Cost + Other_Cost + dist_A1 + (1 + Self_Cost + Other_Cost + dist_A1 | subj)	14	-1269.5	2566.9	6.1^{e-12}	.258	.842
M5_Exp	~ (Self_Cost + Other_Cost + dist_A1)*Exp + (1 + Self_Cost + Other_Cost + dist_A1 | subj)	18	-1239.6	2515.3	**1.00**	**.349**	.838

Table B.34: Results of Model 5, logistic regression predicting the probability of choosing the square object (A1) by Self and Other Disparities and Distance from Cursor in Experiments 1 and 2. Exponentiated estimates are reported with 95% confidence intervals are included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

Model formula: Object Choice ~ Self_Cost + Other_Cost + dist_A1 + (1 + Self_Cost + Other_Cost + dist_A1 | subj)

	Model 5 - Data collapsed over Experiments			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	9.778 (6.156 – 15.531)	0.236	9.657	**<0.001**
Self Disparity	0.995 (0.992 – 0.997)	0.001	-5.031	**<0.001**
Other Disparity	0.993 (0.990 – 0.995)	0.001	-5.903	**<0.001**
Distance from Cursor	0.985 (0.981 – 0.989)	0.002	-7.617	**<0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ \text{subj2}}$	3.43			
$\tau_{11\ \text{subj2.Self\_Cost}}$	0.00			
$\tau_{11\ \text{subj2.Other\_Cost}}$	0.00			
$\tau_{11\ \text{subj2.dist\_A1}}$	0.00			
ρ_{01}	-0.06			
	-0.05			
	-0.71			
ICC	0.79			
N_{subj2}	100			
Observations	3175			
Marginal R^2 / Conditional R^2	0.258 / 0.842			
AIC	2566.947			

Table B.35: Results of Model 5, logistic regression predicting the probability of choosing the square object (A1) by Self and Other Disparities and Distance from Cursor, together with the binary predictor of Experiment ID (1 = Exp.1, 2 = Exp.2) in Experiments 1 and 2. Exponentiated estimates are reported with 95% confidence intervals are included in brackets. Random variances are reported for residuals, individual intercepts and slopes, as well as random-intercept-slope correlations.

Model formula: Object Choice ~ (Self_Cost + Other_Cost + dist_A1) * exp + (1 + Self_Cost + Other_Cost + dist_A1 | subj)

	Model 5, with added Experiment ID - Data collapsed over Experiments			
Predictors	*Odds Ratios*	*std. Error*	*z*	*p*
Intercept	35.395 (19.505 – 64.230)	0.304	11.730	**<0.001**
Self Disparity	0.991 (0.988 – 0.994)	0.002	-5.553	**<0.001**
Other Disparity	0.992 (0.988 – 0.996)	0.002	-4.349	**<0.001**
Distance from Cursor	0.975 (0.970 – 0.980)	0.003	-9.667	**<0.001**
Experiment[2]	0.087 (0.040 – 0.187)	0.390	-6.262	**<0.001**
Self Disparity X Experiment[2]	1.006 (1.002 – 1.010)	0.002	2.862	**0.004**
Other Disparity X Experiment[2]	1.002 (0.997 – 1.007)	0.002	0.671	0.502
Distance from Cursor X Experiment[2]	1.020 (1.013 – 1.027)	0.003	5.606	**<0.001**
Random Effects				
σ^2	3.29			
$\tau_{00\ \mathrm{subj2}}$	1.60			
$\tau_{11\ \mathrm{subj2.Self\_Cost}}$	0.00			
$\tau_{11\ \mathrm{subj2.Other\_Cost}}$	0.00			
$\tau_{11\ \mathrm{subj2.dist\_A1}}$	0.00			
ρ_{01}	0.27			
	0.08			
	-0.56			
ICC	0.75			
N_{subj2}	100			
Observations	3175			
Marginal R^2 / Conditional R^2	0.349 / 0.838			
AIC	2515.292			

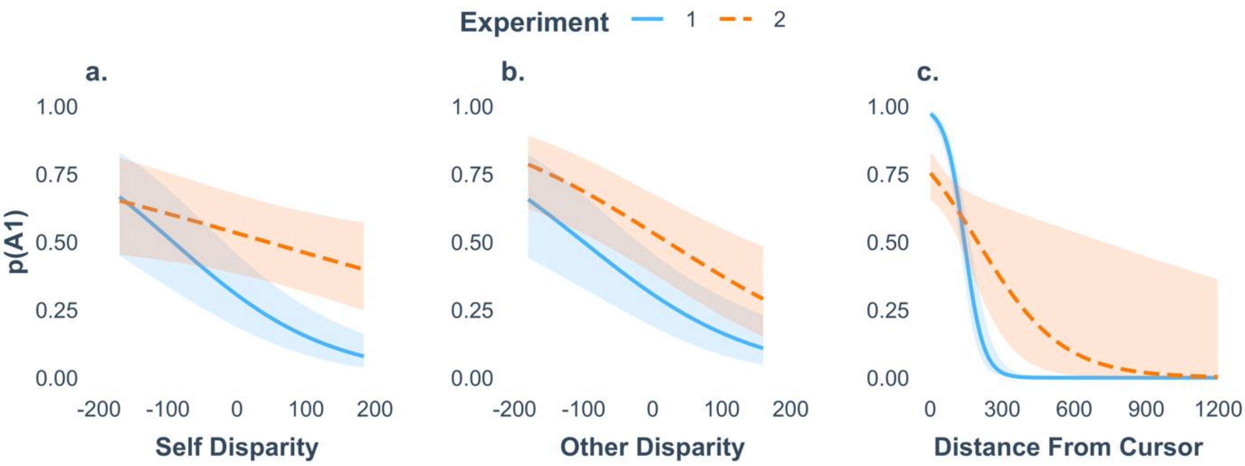

Figure B.7: Predicted A1 choice probabilities in Experiments 1 and 2 (N = 100), according to the best-fitting Model 5 estimated on the pooled data, including the predictor *Experiment*. Decision probabilities are shown as a function of (a) Self Disparity (b) Other Disparity and (c) Distance from Cursor. The shaded ribbons show 95% confidence intervals.

ACKNOWLEDGEMENT

TO EXTERNAL FUNDING AGENCIES

CONTRIBUTING TO THE PHD DISSERTATION

Name of doctoral candidate: Georgina Török

Title of dissertation: *Co-efficiency - An empirical investigation of the principle of rational joint action*

Name of supervisor(s): Natalie Sebanz, Gergely Csibra

External funding agency: European Research Council

Acknowledgement:

This research has received funding from the European Research Council (ERC) under the European Union's Seventh Framework Programme (FP7/2007-2013) under **grant agreement No 616072 – JAXPERTISE.**