
ECONtribute
Discussion Paper No. 407

**Cooperation and Coordination When Others
May Use AI**

Dominik Atella-Suri

Sebastian Kube

May 2026

www.econtribute.de



UNIVERSITÄT
ZU KÖLN

Cooperation and coordination when others may use AI*

Dominik Atella-Suri^{a, b, †}, Sebastian Kube^{b, c}

^a Institute for Food and Resource Economics, University of Bonn, Germany

^b Center for Economics and Neuroscience, University of Bonn, Germany

^c Institute for Applied Microeconomics, University of Bonn, Germany

May 2026

Abstract: Artificial intelligence (AI) is increasingly becoming part of economic decision-making. Yet, in many strategic interactions, individuals may not know whether others rely on AI when forming their decisions. We examine whether decision-makers who are themselves not allowed to use AI behave differently when other group members may consult AI. In an incentivized experiment with a public goods game and a weakest-link game, we exogenously vary whether group members are allowed to use AI to inform their decisions. We find that AI can affect strategic interaction even when it is not directly used by the decision-maker: merely knowing that others may use AI reduces cooperation in the public goods game and effort provision in the weakest-link game. Participants also perceive group members who may use AI as socially more distant and report lower beliefs about appropriate and expected contributions and effort levels. At the same time, the shares of conditional cooperators and conditional coordinators remain largely stable across treatments. These findings suggest that AI is not only a private decision aid but can also shape the social and strategic environment in which economic decisions are made.

JEL classification codes: C71, D83, D91.

Keywords: Cooperation, Coordination, Human-AI Interaction, Artificial Intelligence, Experiment

***Funding:** This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/2 – 390838866. **Acknowledgments:** We thank Lotta Schils and Carl Tamussino for excellent research assistance. We thank Holger Gerhardt for valuable comments. **Research transparency:** Data and analysis code will be available on OSF upon publication. **Authorship contribution statement:** DAS: Conceptualization, Data curation, Formal analysis, Methodology, Visualization, Writing – original draft, Writing – review & editing. SK: Conceptualization, Methodology, Writing – original draft, Writing – review & editing. **Declaration of AI and AI-assisted technologies in the writing process:** During the preparation of this work, the authors used ChatGPT in order to improve readability and language. After using this service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication. **Declaration of competing interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

[†]DAS is the corresponding author. Address: Nussallee 21, 53115 Bonn, Germany. E-mail: dsuri@uni-bonn.de.

1 Introduction

Artificial intelligence (AI)—technologies that can execute tasks one may expect to require human intelligence (Kaplan, 2016)—is rapidly becoming part of economic decision-making. Large language models (LLMs) and other AI systems are used to generate advice, summarize information, evaluate options, and support decisions in a wide range of contexts (e.g., Tse et al., 2024; Ben-Michael et al., 2025; Gieselmann et al., 2025; Huseynov, 2025; Yin et al., 2025). This development has generated a rapidly growing literature on the behavioral consequences of AI (Chugunova & Sele, 2022). Existing work has studied, among other things, whether people follow AI-generated advice and if this changes ethical behavior (Köbis et al., 2021; Leib et al., 2025; Greiner et al., 2025), how people evaluate decisions made by AI rather than humans (Leib et al., 2024; Wester et al., 2024; Chugunova & Luhan, 2025), or how (in particular generative) AI may transform occupations and labor-market prospects (Noy & Zhang, 2023; Orchard & Tasiemski, 2023; Brynjolfsson et al., 2025; Teutloff et al., 2025). A related stream of research examines AI and large language models as social actors in strategic environments, asking how AI agents behave in economic games and how humans respond in strategic interactions with AI (Karpus et al., 2021, 2025; Mei et al., 2024; March, 2021; Akata et al., 2025; Dvorak et al., 2025; Xie et al., 2025).

Much of this work focuses on direct exposure to AI: individuals receive advice from AI, delegate a decision to AI, interact with an artificial agent, or evaluate an outcome that is known to have been generated by AI. Yet many economically relevant situations involve a more indirect form of AI exposure. For example, in workplaces, markets, negotiations, online platforms, and group decision-making environments, individuals may not use AI themselves but may know that others could do so. A worker may not consult a LLM when preparing a proposal but may expect that colleagues or competitors do. A negotiator may rely on her own judgment but believe that the other side may have used AI support. A team member may make an effort choice without AI assistance while knowing that other team members might have received AI-generated advice. In such environments, AI becomes part of the strategic background: it may affect behavior not because it directly changes one's own decision process, but because it changes beliefs about the reasoning, motives, and likely behavior of others.

Our paper studies this indirect channel. Specifically, we ask whether individuals change their behavior when they know that other group members may use AI to support their decisions, even though they themselves are not allowed to use AI. This question is important because many consequences of AI adoption may arise beyond direct AI use. The diffusion of AI may alter expectations about other people's behavior, perceptions of social norms, and the strategic uncertainty surrounding human interaction. If individuals believe that AI-assisted others are more selfish, more strategic, more sophisticated, or less socially connected, they may respond differently even when all final decisions are still made by humans.

We investigate this question experimentally using two major classes of interaction—cooperation and coordination—which capture distinct but central problems in economics. To study cooperation, we use a one-shot public goods game (PGG). Here, the main problem is free-riding: a group is better off if everyone contributes, but each individual has a selfish monetary incentive not to contribute while benefiting from others' contributions. For players that are not self-centered money maximizers, behavior depends (not only) on perceived

norms and expectations about others' contributions (Gächter & Thöni, 2005; Chaudhuri et al., 2017; Angelovski et al., 2018; Kölle & Quercia, 2021). To study coordination, we use a one-shot weakest-link game (WLG): high effort pays off when others also choose high effort, but is costly when even one group member chooses a lower effort level. While each individual's best response is to match the expected group minimum, not knowing what effort others will do introduces strategic uncertainty as the core driver of coordination failure (Huyck et al., 1990; Devetag & Ortmann, 2007; Chaudhuri et al., 2015; Bortolotti et al., 2016).

The comparison between the PGG and the WLG allows us to test whether the possibility of AI use by others has domain-specific indirect effects. Our a priori reasoning was that participants expect AI advice to emphasize rationality and strategic self-interest—in which case the anticipation of other group members using AI assistance may have opposite implications for cooperation and coordination. In the PGG, knowing that others may consult AI could hamper cooperation because strategic self-interest would imply free-riding. Under this view, the possibility that some people in a group use AI may make groups less successful: it reduces cooperation not only for those that use AI, but also for those who do not use AI. In the WLG, the same perception of AI as selfish, strategic and rational may have different consequences. People using AI might now better understand the equilibrium structure of the game. Anticipating this may increase trust in the group's ability to reach the payoff-dominant outcome and reduce pessimism about the minimum effort. In that case, the possibility that some people in a group use AI may foster coordination and make groups more successful: it increases effort levels not only for those that use AI, but also for those who do not use AI.

To test these ideas, we conducted a preregistered online experiment with participants from the BonnEconLab subject pool at the University of Bonn. Participants were randomly assigned to one of three treatment conditions. Treatments differed in whether participants themselves were allowed to use AI and whether others in their group were allowed to do so. In the NoAI-NoAI condition, no participant was allowed to use AI and this was common knowledge. In the NoAI-AI condition, participants were not allowed to use AI but were informed that this restriction did not apply to all group members, so they could expect that at least some group members would use AI to support their decisions. In the AI-AI condition, participants were allowed to use AI themselves and were informed that some other group members might also use AI. Our main causal comparison is between the NoAI-NoAI and NoAI-AI conditions. In both conditions, the focal participant is prohibited from using AI. The only difference is whether the other group members may have access to AI. This comparison, varying others' possible AI use while holding constant one's own AI access, thus allows us to causally identify the effects of the indirect channel that AI might have on behavior in the PGG and the WLG.

For each game, we elicited both an unconditional decision and a conditional decision profile in an incentive-compatible way. Following Fischbacher et al. (2001), participants in the PGG indicated how much they would contribute conditional on the average contribution of the other three group members; resp. how much effort to exert conditional on the the other group members' minimum effort in the WLG. This allows us to study not only whether anticipated AI use affects contribution and effort levels per se, but also whether it changes the structure of how people respond to others' behavior. In addition, we asked what people believe to be the socially appropriate thing to do in each game (first-order normative beliefs), elicited what they think others believe to be the socially appropriate thing to do (second-order

normative beliefs), and what they expect others to actually do (descriptive beliefs).

We find that others' possible AI access has negative indirect effects on both cooperation and coordination. Participants in NoAI-AI contributed significantly less to the public good than participants in NoAI-NoAI (7.39 vs. 9.53 tokens, a reduction of approximately 29%), and chose significantly lower effort levels in the weakest-link game (10.4 vs. 12.62 hours, a reduction of approximately 21%). At the same time, the shares of conditional cooperators and conditional coordinators remain largely stable across treatments. The possibility that others may use AI is also reflected in participants' beliefs and social perceptions: in the PGG, second-order normative beliefs are significantly lower in NoAI-AI than in NoAI-NoAI, while in the WLK, first-order normative beliefs and descriptive beliefs are significantly lower. Moreover, they perceive group members who may use AI as socially more distant than group members who cannot.

Our study contributes to three strands of literature. The first is the extensive literature exploring cooperation, in particular with respect to the voluntary contributions to the provision of public goods. A well-established finding in public goods experiments is that voluntary cooperation is fragile and often declines over time, with free riding eventually becoming prevalent (Zelmer, 2003; Neugebauer et al., 2009; Fischbacher & Gächter, 2010; Jouxtel, 2019; Catola et al., 2023). Prominent solutions to the cooperation problem are often built around punishment (e.g., Choi & Ahn, 2013; Goeschl & Jarke, 2016; Grieco et al., 2017; Tan & Xiao, 2018; Morgan et al., 2019) or related channels that (either implicitly or explicitly) alter the underlying incentive structure (e.g., Drouvelis et al., 2017) to make cooperation more appealing than free-riding in relative terms (see Chaudhuri, 2011, for a survey). Another approach focuses on information provision (see Engel et al., 2021, and the references therein), as this might alter empirical and normative expectations which in turn are likely to affect cooperation (e.g., Guillen et al., 2010; Guala et al., 2013; Gächter et al., 2017; Zhang, 2019).

Our study complements these approaches by showing that the tools and sources of advice available to other group members might constitute relevant information that shapes (beliefs and) cooperation, since we find that unconditional cooperation is lower in NoAI-AI than in NoAI-NoAI. Moreover, our findings underline the robustness of fundamental cooperation types, as these have been observed in periods when general-purpose AIs were not publicly available (Volk et al., 2012; Thöni & Volk, 2018). This is also noteworthy given recent evidence by Engl et al. (2021) indicating that cooperation preferences might be shaped by the available institutions, under which the presence of advice-giving tools (such as modern AI) may also fall.

A second strand of literature to which our study contributes examines coordination under strategic uncertainty, in particular in weakest-link and minimum-effort games. A large experimental literature has shown that groups often fail to coordinate on payoff-dominant equilibria, especially when successful coordination requires all group members to choose high effort and when subjects are uncertain about others' willingness to do so (Huyck et al., 1990; Devetag & Ortmann, 2007). At the same time, this literature also shows that coordination failure is not inevitable. Efficient coordination can be fostered by mechanisms that create common expectations, provide focal points, or reduce uncertainty about others' behavior, such as pre-play communication (Blume & Ortmann, 2007), managerial communication (Brandts & Cooper, 2007), changes in incentives that serve as triggers for coordinated change (Brandts

& Cooper, 2006), gradual group growth with observable successful histories (Weber, 2006), or information generated by related market or organizational environments (Kogan et al., 2011; Chaudhuri et al., 2015; Bortolotti et al., 2016). Recent evidence also highlights the importance of social relationships in weakest-link team production: Gächter et al. (2025) show that group cohesion is strongly associated with performance in weak-link games, with high-cohesion groups being more likely to achieve superior equilibria, and that this association appears to operate primarily through beliefs rather than social preferences.

Our study adds to this literature by considering a novel source of strategically relevant information: whether other group members may rely on AI-based advice when forming their effort decisions. Unlike communication, bonuses, or shared histories, others' possible AI access does not provide a common recommendation, does not change the material payoff structure, and does not necessarily create a focal point for high-effort coordination. Instead, it may change beliefs about how others reason, how predictable their choices are, and whether they will follow payoff-maximizing strategies. Contrary to our initial expectation that AI access by others would reduce coordination risk and thereby increase effort provision, we find that unconditional effort is lower in NoAI-AI than in NoAI-NoAI. Thus, in our setting, the possibility that others may use AI appears to make high-effort coordination less attractive rather than more attainable. At the same time, the conditional effort profiles remain stable, suggesting that established types of conditional coordination are robust even when generative AI becomes part of the strategic background.

A third strand of literature studies how AI enters strategic interaction. This literature can be divided into two related approaches. One approach varies the identity of the interaction partner and asks how humans behave when the counterpart is an artificial agent rather than another human. Evidence from social dilemmas suggests that people may cooperate less with artificial agents than with humans (Karpus et al., 2021). Relatedly, Ishowo-Oloko et al. (2019) show in a repeated prisoner's dilemma that bots can elicit more cooperation than humans, but that this advantage disappears when participants are informed that they are interacting with a bot. Other work shows that machines can sustain cooperation with humans in repeated games when equipped with appropriate signaling mechanisms (Crandall et al., 2018), and that LLMs themselves display distinctive patterns of cooperation and coordination in repeated games (Akata et al., 2025). Recent AI-based studies further show that humans adjust their strategic choices when facing AI opponents (Barak & Costa-Gomes, 2025; Jiang et al., 2026). A related set of studies examines mixed human-AI groups. In a networked coordination game, Shirado and Christakis (2017) show that centrally placed bots with moderate behavioral noise can improve global human coordination. On the other hand, Dell'Acqua et al. (2025) show in a team-production experiment that introducing automated AI co-workers can reduce team performance, increase coordination failures, and lower trust and effort provision, even when the automated agent outperforms human workers.

A second approach, closer to ours, treats AI as a tool or decision aid that changes the strategic environment among human decision-makers. Dvorak et al. (2025) examine strategic interactions in which AI enters the decision-making process of one side. In contrast to our setting, however, AI directly takes over the partner's decision across several two-player games. They find that fairness, trust, trustworthiness, cooperation, and coordination decrease, but only when AI's role in the partner's decision is disclosed; not when participants are uncertain about

whether they face a human or an AI. Most closely related to our design, Shen et al. (2025) study asymmetric AI assistance in a trust game and a prisoner’s dilemma. They find that, although AI assistance embedded as an LLM-chatbot in the decision interface does not significantly change overall cooperation rates, non-assisted players perceive AI-assisted counterparts as more competitive and less “warm”, where lower warmth refers to lower perceived benevolence, prosociality, or cooperative intent. Our study complements this evidence by focusing on a more indirect channel: the focal participant is not allowed to use AI, the final decisions are still made by human group members, and AI may or may not actually be used by others. Moreover, we study both cooperation and coordination, and elicit decisions both directly and with the strategy method (Selten, 1967). This not only allows us to isolate whether anticipated AI-assisted decision-making by others changes cooperation and coordination, but also whether such changes are accompanied by shifts in general cooperation/coordination patterns, beliefs, or perceived social distance.

The remainder of the paper is organized as follows. In the next section, we will explain our experimental design and procedures and motivate our preregistered behavioral hypotheses. We then move on to show that randomization into treatments was successful w.r.t. observables, before showing the results from our two main treatments for both the cooperation and coordination game. We end with a discussion of our results, our design and potential next steps for future research.

2 Material and methods

In this section, we first describe the details of the two games that we used, the PGG and the WLG. Afterwards, we describe the treatments and the preregistered hypotheses. In the end, we provide details on the experimental procedure as well as on the sample characteristics.

2.1 Public goods game (PGG)

We implemented a standard linear public goods game using parameters that are equal or very similar to what has been used in the corresponding literature (e.g., Fischbacher et al., 2001, 2012; Kocher et al., 2015; Grieco et al., 2017; Zhang, 2019). Each participant $i \in \{1, \dots, 4\}$ was endowed with 20 tokens, which could be allocated either to a public account, denoted by g_i , or to a private account ($20 - g_i$). Each token allocated to the public account generated a marginal return of 0.4 coins for each group member, where one coin corresponded to 0.1 Euro. By contrast, tokens allocated to the private account yielded a return only to the individual participant, at a rate of one token per coin. Accordingly, participant i ’s payoff is given by $\pi_i = 20 - g_i + 0.4 \sum_{n=1}^4 g_n$.

Under the standard assumption of self-interested payoff maximization, the unique subgame-perfect Nash equilibrium is zero contribution to the public account. By contrast, the social optimum is attained when all participants contribute the maximum amount, thereby maximizing total group welfare.

2.2 Weakest-link game (WLG)

The weakest-link game was implemented as follows, in line with previous studies (e.g., Albrecht & Kube, 2018; Gächter et al., 2025). Each participant $i \in \{1, \dots, 4\}$ was endowed with 20

hours representing effort levels, which could be allocated either to a public project, denoted by g_i , or to a private project ($20 - g_i$). Each hour allocated to the private project yielded a return of one coin to the individual participant only. In the public project, payoffs were determined solely by the smallest effort made by any group member. Specifically, each group member received 1.6 coins for every hour contributed at the minimum effort level within the group. Accordingly, participant i 's payoff is given by $\pi_i = 20 - g_i + 1.6 \times \min_n(g_n)$ with $n \in \{i, j, k, l\}$.

Any common effort level chosen by all group members constitutes a Nash equilibrium, i.e., any profile satisfying $g_i = g_j = g_k = g_l$. These equilibria can be ranked by efficiency. In particular, the equilibrium with $g_i = g_j = g_k = g_l = 20$ is the most efficient and payoff-dominant Nash equilibrium, whereas the equilibrium with $g_i = g_j = g_k = g_l = 0$ is the least efficient but risk-dominant equilibrium.

2.3 Treatment conditions and hypotheses

Our experiment varies three dimensions: i) the type of game participants play, ii) whether participants themselves are allowed to use AI, and iii) what participants know about the possible AI use of the other group members. The first dimension is implemented within subjects: all participants first played the public goods game (PGG) and then the weakest-link game (WLG). Before making their decisions in each game, participants were informed whether they themselves were allowed to use AI and whether other group members might be allowed to use AI.

With respect to participants' own AI access, we distinguish between participants with AI access and participants without AI access. Participants with AI access were told: "You are allowed to use or consult artificial intelligence to make your decisions." Participants without AI access were told: "You are not allowed to use or consult artificial intelligence to make your decisions."

All participants with AI access received identical information about the other group members. They were informed that "Not every member was explicitly allowed to use artificial intelligence to make decisions. Therefore, it may be that some of your group members make their decisions without the help of artificial intelligence." We refer to this treatment as AI-AI. Since our main interest lies in the indirect effect of others' possible AI use on participants who do not use AI themselves, results for this treatment are reported in the Online Appendix. The treatment was included to implement the two main treatment conditions without deception.

Our two main treatment conditions are NoAI-NoAI and NoAI-AI. Both involve participants who were not allowed to use AI themselves. In the NoAI-NoAI treatment, participants were additionally informed that "The other members of your group have received the same message." Thus, neither the focal participant nor the other group members were allowed to use AI. In the NoAI-AI treatment, participants were instead informed that "Not every member was explicitly prohibited from using artificial intelligence to make decisions. Therefore, it may be that some of your group members make their decisions with the help of artificial intelligence." Thus, participants themselves were not allowed to use AI, but other group members might have been allowed to do so.

Taken together, participants made decisions in one of three environments. In AI-AI, participants were allowed to use AI and knew that some other group members might also use AI. In NoAI-NoAI, participants were not allowed to use AI and knew that the same restriction applied to all other group members. In NoAI-AI, participants were not allowed to use AI but

knew that some other group members might use AI. Our main behavioral predictions concern the comparison between NoAI–NoAI and NoAI–AI. In both conditions, the focal participant is prohibited from using AI. The only difference is whether other group members may have access to AI. This comparison therefore identifies the causal effect of others' possible AI access while holding participants' own AI access constant.

For this comparison, we preregistered the following hypotheses:

Hypothesis 1: Cooperation in a public goods game decreases if other group members may use AI to form their contribution decision.

Hypothesis 2: Coordination in a weakest-link game increases if other group members may use AI to form their effort decision.

The rationale behind our hypotheses rests on the idea that participants may associate AI-supported decision-making with more strategic, payoff-oriented, and instrumentally rational reasoning. Several findings are consistent with this view. In a review of experiments involving computer players, March (2021) concludes that humans often adapt their behavior when interacting with artificial agents and tend to act more selfishly and more rationally in their presence. Related, Barak and Costa-Gomes (2025) show that participants choose more equilibrium-like actions when facing LLM opponents in a strategic game, suggesting that humans may expect LLMs to reason more strategically than humans. Evidence on AI advice points in a similar direction: Leib et al. (2024) show that AI-generated advice can affect ethically relevant behavior, while Jobu Babin and Chauhan (2026) find that generative-AI advice in public-good and coordination-type games emphasizes guided reasoning and concepts of risk and strategy. Finally, Shen et al. (2025) show that non-assisted participants perceive AI-assisted counterparts as more competitive and less warm, i.e., as less benevolent or less prosocially oriented.

Taken together, this evidence suggests that participants in our NoAI–AI condition may expect AI-assisted group members to reason more strategically and less prosocially. In the public goods game, such expectations point toward lower cooperation, because payoff-maximizing behavior implies free-riding. In the weakest-link game, however, the same expectation may imply higher effort provision, because strategically sophisticated group members may be expected to recognize the payoff-dominant equilibrium and to coordinate on high effort. We therefore hypothesized that others' possible AI access decreases cooperation in the public goods game, but increases coordination in the weakest-link game.

At the same time, one should note that the direction of these effects is theoretically ambiguous, since alternative predictions might also be derived from the literature. First, AI advice does not need to be selfish. Jobu Babin and Chauhan (2026) find that both human and generative-AI advice increase prosociality, including contributions in a public goods game, although human advice appears more norm-oriented and more effective. Similarly, Capraro et al. (2025) show that advanced LLMs can capture qualitative patterns in other-regarding behavior but tend to underestimate self-interest and overestimate altruistic behavior. Second, people sometimes expect AI systems to make more utilitarian or outcome-oriented decisions than humans (Zhang et al., 2022b; Myers & Everett, 2025). In the public goods game, such utilitarian reasoning could imply higher rather than lower contributions, since full contribution maximizes total surplus. Third, current LLMs may (still) struggle with complex multi-

agent reasoning and socially adept coordination under uncertainty (Yao et al., 2025). In the weakest-link game, this could imply that others' possible AI access increases rather than decreases strategic uncertainty, especially if participants expect different group members to receive different or poorly coordinated advice. Thus, while our preregistered hypotheses reflect the view that AI access makes others appear more strategic and payoff-oriented, the literature also provides reasons to expect higher cooperation, lower coordination, or no systematic effect.

As an auxiliary check on the kind of advice participants might plausibly expect from LLMs, we simulated the advice that contemporary models would provide in our experimental setting. Pretending to be a participant in our study, we entered the relevant decision problems into ChatGPT 5.3 Instant and Gemini 3 Fast a total of 200 times, with 100 independent prompts per model. The results reveal substantial heterogeneity across models. In the one-shot public goods game, ChatGPT 5.3 Instant recommended contributing 0 tokens in 95 out of 100 runs. Gemini 3 Fast, by contrast, recommended contributing 0 tokens in only 43 out of 100 runs and recommended contributing the full endowment of 20 tokens in 54 out of 100 runs. In 10 out of 100 runs, Gemini explicitly highlighted the trade-off between the individually strategic choice and the socially efficient outcome. In the one-shot weakest-link game, ChatGPT 5.3 Instant recommended choosing 0 hours in 81 out of 100 runs and recommended choosing the maximum effort of 20 hours in only 11 out of 100 runs. Gemini 3 Fast, by contrast, recommended the maximum effort of 20 hours in 95 out of 100 runs.

The conditional decision tasks show similarly LLM-specific patterns. In the conditional public goods game, ChatGPT 5.3 Instant recommended a free-riding profile in 86% of cases, compared to 43% for Gemini 3 Fast. Gemini 3 Fast recommended matching the average contribution of the other group members in 54% of cases, consistent with a conditional-cooperation logic. In the conditional weakest-link game, ChatGPT 5.3 Instant recommended choosing 0 hours irrespective of the minimum effort of the other group members in 53% of cases and recommended matching the weakest link in 45% of cases. Gemini 3 Fast recommended matching the weakest link in all cases. These simulations illustrate two points. First, LLM advice can be strategically consequential in both games, as the recommendations map directly onto contribution and effort choices. Second, the direction of advice is highly model-dependent: while ChatGPT 5.3 Instant predominantly recommended payoff-maximizing low contributions and low effort, Gemini 3 Fast more often recommended socially efficient or coordination-oriented choices. Thus, participants' beliefs about others' possible AI use could reasonably support different behavioral responses, making the effect of anticipated AI assistance an empirical question.

2.4 Procedures

Our study received ethical approval from the German Association for Experimental Economic Research e.V. (No. jd2mHZ4m) and was preregistered on AsPredicted.org (<https://aspredicted.org/fb3h85.pdf>). The experiment was hosted on Qualtrics LLC. We conducted the study online drawing on the BonnEconLab subject pool at the University of Bonn (managed using hroot; Bock et al., 2014). In total, we invited $N = 241$ participants. Using Qualtrics' built-in randomization function, participants were randomly assigned to one of three treatment conditions (NoAI-NoAI, NoAI-AI, and AI-AI; see Section 2.3), with a target of 80 participants per

condition. The sample size was determined based on an a priori one-sided power analysis with $\alpha = .05$ and 80% power to detect a medium effect size of $d = 0.4$.

After clicking on the web link to access our study, participants were required to provide informed consent; otherwise, participation was not possible. They were then presented with the description of the public goods game (see Section 2.1) and asked to make an incentivized unconditional contribution decision to the public good. After making their unconditional allocation decision between the private and the public good, we used a variant of the strategy-method approach (Selten, 1967) and asked participants to make the same decision again conditional on the average contribution of the other three group members to the public good. Following Fischbacher et al. (2001), one player per group was randomly selected *ex post* to serve as the conditional decision-maker, whereas the remaining players served as unconditional decision-makers. Since at the time a player makes the unconditional and the conditional decisions she does not know yet which of the decisions will become payoff-relevant for her, incentive compatibility is granted for all decisions.

Next, participants reported their normative and descriptive beliefs (Bicchieri, 2006) about others' contributions to the public good. For this, we first elicited first-order normative beliefs about socially appropriate contribution levels to the public good. Then, we asked participants for their second-order normative beliefs, i.e., their beliefs about what others would consider socially appropriate. This captures perceived social norms of cooperation. We incentivized the elicitation of the second-order normative belief using the coordination approach of Krupka and Weber (2013), i.e., using other players' answer to the first-order normative beliefs to incentivize the belief about others' second-order normative belief. To measure descriptive beliefs, we asked participants to report the average contribution they expected from others. The accuracy of both second-order normative beliefs and descriptive beliefs was incentivized using the quadratic scoring rule (Selten, 1998).

After the public goods game, participants played the weakest-link game (see Section 2.2) with the same group members. The procedure in this game exactly mirrored that of the public goods game: Participants first made an unconditional effort decision, followed by conditional effort choices and belief elicitation.

The experiment concluded with a questionnaire. First, we elicited subjective closeness to group members using the "Inclusion of the Other in the Self" scale (IOS11; Baader et al., 2024). In the NoAI-AI and AI-AI treatments, participants were specifically asked about their closeness to group members who were allowed to use AI. We also administered the ATTARI-12 questionnaire to measure attitudes toward artificial intelligence (Stein et al., 2024). For the ATTARI-12 index, items 2, 4, 7, 8, 10, and 12 were reverse-coded, and the mean of the 12 items, each measured on a 5-point Likert scale, was calculated. Following the ATTARI-12 questionnaire, we elicited trust in AI using the five items reported in Suri et al. (2026a), each measured on a 5-point Likert scale. We also followed their approach to construct the trust index. In addition, we measured general trust in humans and positive reciprocity using the self-assessment items from the Global Preference Survey (Falk et al., 2018). Finally, we elicited participants' gender and age.

Data collection took place in April 2026. We stopped data collection once the target of 80 participants per treatment condition was met. Due to the randomizer of Qualtrics, we ended up with 81 participants in the condition NoAI-NoAI and with 80 participants in the other two

conditions. The experiment was administered in German. Screenshots will be displayed in the OSF repository. English translations are presented in the Online Appendix. The target completion time was 20 minutes. Participants received a fixed participation payment of 4 Euro and could earn additional compensation based on their own and their group members' decisions in the games. For this, we first randomly matched participants into groups of 4 within the NoAI-NoAI treatment separately and the NoAI-AI and AI-AI treatments jointly. For each group, we then randomly selected one of the two games (i.e., PGG or WLG), followed by another random draw on whose group member's conditional decision is realized. This procedure was known to all participants, though they did not know *ex ante* which game would be randomly selected to ensure incentive-compatibility.

2.5 Our participant pool

We stopped data collection after reaching our preregistered target sample size of 80 participants per treatment condition. Because Qualtrics randomly assigned participants to conditions, the final sample comprised 81 participants in NoAI-NoAI and 80 participants in each of the other two conditions. Table 1 reports descriptive sample characteristics for the full sample and separately by treatment condition. Overall, 46.8% of participants identify as male, and the average age is 28.9 years. Gender and age do not differ significantly across treatment conditions. Similarly, we find no significant differences across conditions in trust in other humans, trust in AI (Cronbach's $\alpha = .77$), attitudes toward AI (measured by ATTARI-12, Cronbach's $\alpha = .91$), or positive reciprocity—indicating successful randomization into treatment conditions. Note that the two trust measure scores are not directly comparable, as trust in AI is measured on a 5-point Likert scale, whereas trust in humans is measured on an 11-point Likert scale.

Table 1: Sample characteristics

Characteristic	Full	NoAI-NoAI	NoAI-AI	AI-AI	p-value
Participants	241	81	80	80	
Male	46.8%	51.9%	41.6%	46.8%	.432
Age	28.9	29.4	29.9	27.5	.285
Positive Reciprocity	8.7	8.8	8.5	8.8	.223
Trust in Human	4.6	4.5	4.7	4.7	.886
Trust in AI	2.5	2.4	2.6	2.5	.072
ATTARI-12	3.1	3.1	3.2	3.1	.569

Notes: The table reports mean value sample characteristics with regard to the full sample and separately for each treatment condition. The p-values for the comparison among all three treatment conditions originates from Pearson's Chi-squared or Kruskal-Wallis rank sum tests, depending on the nature of the characteristic.

3 Results

We present the results by focusing on the two main treatment conditions that are directly relevant to our preregistered hypotheses: NoAI-NoAI and NoAI-AI. Results for the AI-AI condition, as well as additional tables and figures, are reported in the Online Appendix for completeness and exploratory interpretation. All analyses were conducted in R, version 4.2.2 (R Core Team, 2022). Unless otherwise stated, tests of the preregistered directional hypotheses are one-sided, whereas exploratory comparisons are reported using two-sided tests. A replication package will be made available on OSF.

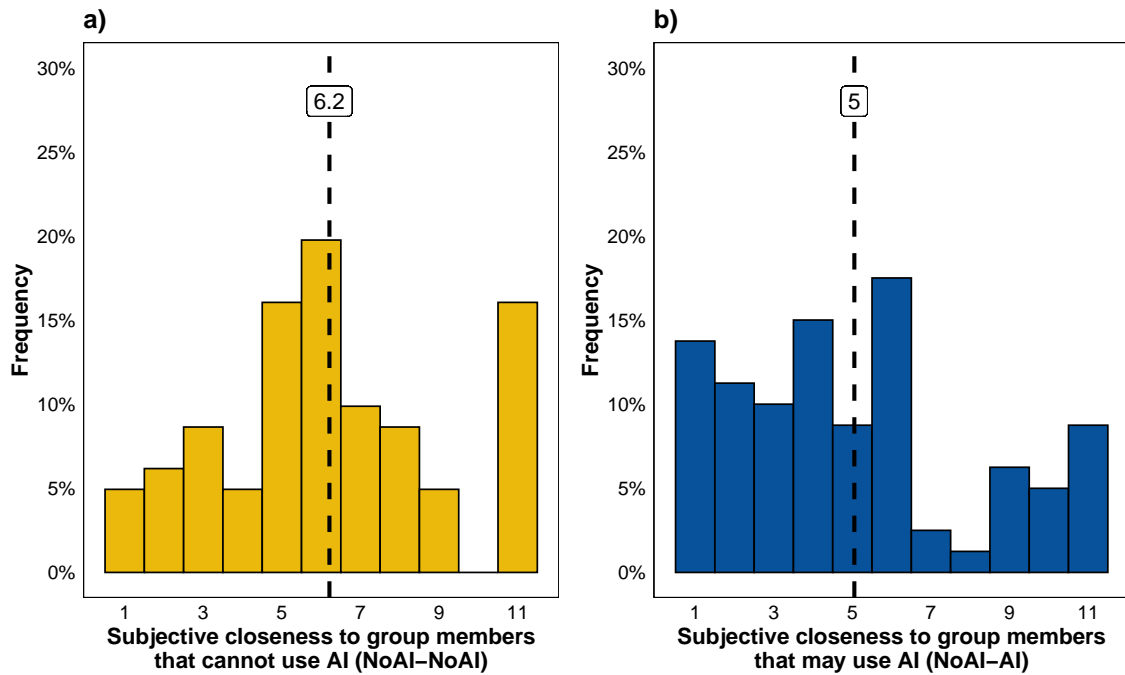


Figure 1: Subjective closeness to group members. The figure shows the distribution of subjective closeness toward other group members that cannot use AI to inform their decisions (left, a, NoAI-NoAI) and other group members that may use AI to inform their decisions (right, b, NoAI-AI). Subjective closeness is measured by the IOS11 score (Baader et al., 2024), where 1 means subjectively very distant and 11 means subjectively very close. The dashed lines represent the respective means.

We first examine participants' perceived subjective closeness to other group members. Figure 1 displays the distribution of IOS11 scores separately for the two main treatment conditions. Panel a shows that in the NoAI-NoAI condition, where participants rated their closeness to other group members who also could not use AI, the distribution is bimodal, with peaks at the midpoint of the scale and at the highest value, indicating maximum perceived closeness. By contrast, in the NoAI-AI condition, where participants themselves were not allowed to use AI but rated their closeness to group members who may use AI, the distribution shifts toward lower values: most reported scores are 6 or below. The two distributions differ significantly (two-sided Kolmogorov-Smirnov test, $p = .003$). The average IOS11 score is also lower in NoAI-AI than in NoAI-NoAI, and a two-sided Wilcoxon rank-sum test indicates a significant difference between conditions ($p = .007$). These findings suggest that participants who are not allowed to use AI perceive themselves as socially more distant from others who may use AI to inform their decision-making. This is interesting, given that previous studies

Table 2: Treatment effects

Variable	Public Goods Game			Weakest-Link Game		
	NoAI-NoAI	NoAI-AI	p	NoAI-NoAI	NoAI-AI	p
Unconditional contribution/ effort	9.53 (7.68)	7.39 (7.01)	.042	12.62 (7.52)	10.4 (6.74)	.031
Conditional cooperator/ coordinator	.73 (.45)	.80 (.40)	.377	.85 (.36)	.86 (.35)	1.00
Free-rider	.12 (.33)	.13 (.33)	1.00	–	–	–
Other	.15 (.36)	.07 (.27)	.222	.15 (.36)	.14 (.35)	1.00
First-order normative belief	7.26 (5.10)	6.47 (4.65)	.332	10.98 (6.36)	8.85 (5.82)	.024
Second-order normative belief	8.64 (4.06)	6.61 (3.90)	<.001	10.05 (5.39)	9.03 (5.24)	.250
Descriptive belief	7.79 (5.35)	7.28 (5.02)	.398	11.38 (5.99)	9.41 (5.93)	.026

Notes: The table reports mean (standard deviation) values for contribution/ effort decisions, conditional types and beliefs. The p-values originate from Pearson's Chi-squared or Wilcoxon rank sum tests—depending on the nature of the characteristic.

have found subjective closeness to matter in both strategic (e.g., Gächter et al., 2022, 2025) and non-strategic (e.g., Suri et al., 2026a, 2026b) decision-making settings.

Table 2 summarizes the main treatment comparisons for the public goods game (PGG) and the weakest-link game (WLG). In the PGG, unconditional contributions to the public good—i.e., contributions made without knowing other group members' decisions—are significantly higher in NoAI-NoAI than in NoAI-AI. The difference amounts to 2.14 contribution points, corresponding to approximately 29% higher contributions in NoAI-NoAI relative to NoAI-AI. Contrary to our expectation where we hypothesized that AI could boost coordination as the reasoning underlying both higher personal and higher group payoffs is aligned, we observe a similar pattern in the WLG. Effort is significantly higher in NoAI-NoAI, corresponding to approximately 21.3% higher effort relative to NoAI-AI. Overall unconditional levels are higher in the WLG than in the PGG, consistent with the different strategic structure of the coordination game.

Next, we classify participants into cooperation and coordination types, following standard approaches in the literature (see, e.g., Fischbacher et al., 2001; Rustagi et al., 2010; Gächter et al., 2017). The resulting patterns are displayed in Figure 2; though we now focus only on free-rider and conditional cooperator types as done in previous studies (see, e.g., Kocher et al., 2015). In the PGG, we define free-riders as participants who never contribute to the public good, irrespective of the contributions of the other group members. The share of free-riders is similar across the two treatment conditions, amounting to 12% in NoAI-NoAI and 13% in NoAI-AI.

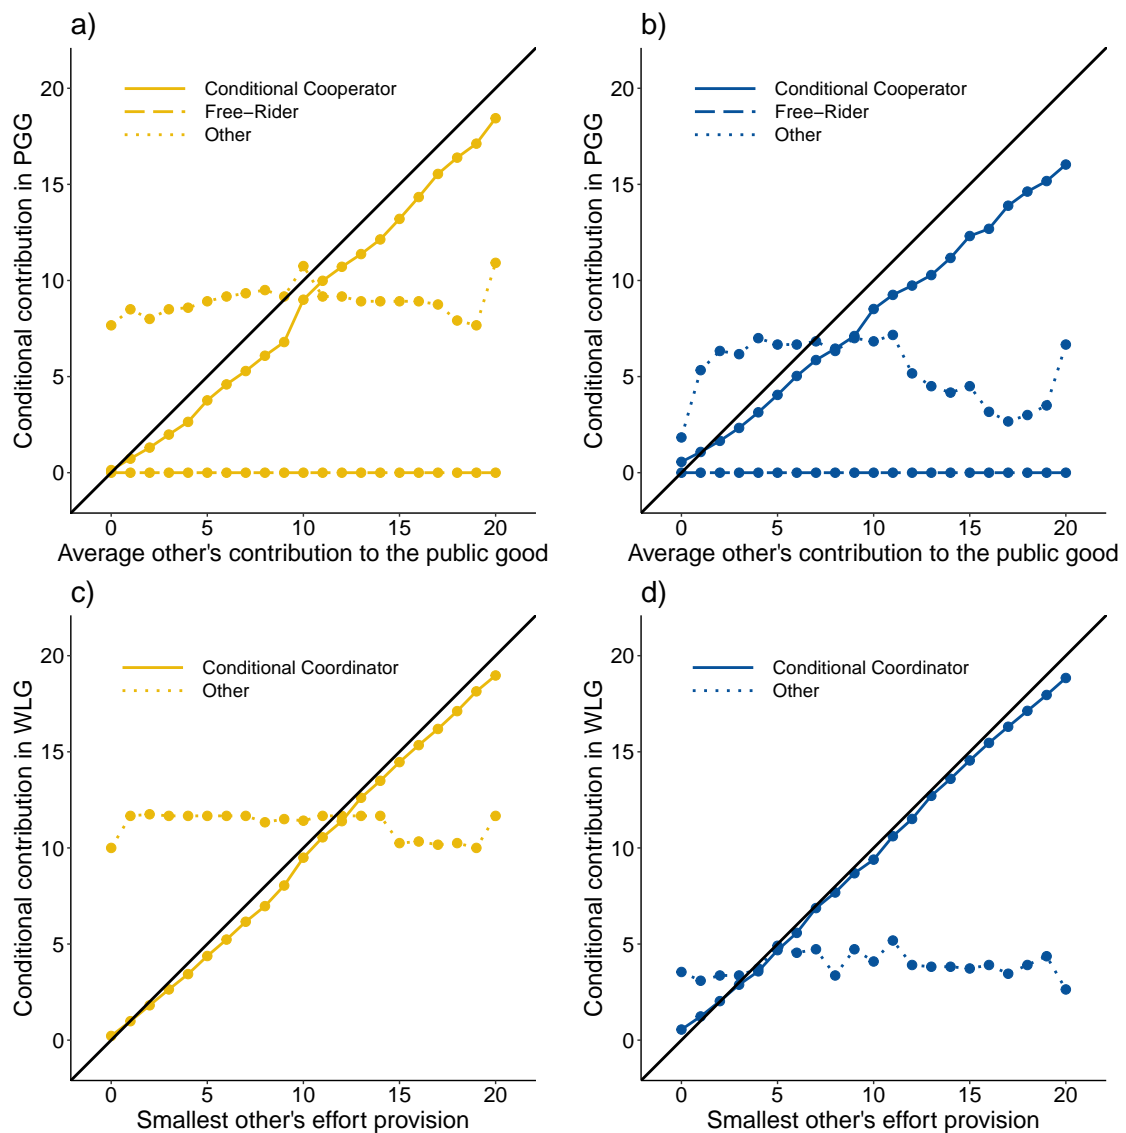


Figure 2: Conditional types. The figure shows the conditional types in the NoAI-NoAI (left) and in the NoAI-AI condition (right). The top row shows the types in the PGG and the bottom row shows the types in the WLG. The solid lines represent the 45 degree line.

We define conditional cooperators as participants whose own contributions increase significantly with the contributions of the other group members, based on a positive Spearman rank correlation that is statistically significant at the 5% level. According to this classification, 73% of participants in NoAI-NoAI and 80% in NoAI-AI are conditional cooperators. This difference is not statistically significant. Although the aggregate conditional cooperation patterns appear broadly similar across treatment conditions, as shown in Panels a and b of Figure 2, the estimated response slopes differ significantly. The slope in NoAI-AI is less steep than in NoAI-NoAI ($\beta = -.14$, significant at the 1% level; see the random-effects estimation in Column 6 of Table OA-1). We also find a significant level difference between conditions ($\beta = .75$, significant at the 5% level; see Column 6 of Table OA-1). This suggests that even among participants who remain conditional cooperators in type, the perceived AI environment dampens the intensity of reciprocal cooperation: conditional cooperators in NoAI-AI respond less strongly to others' contributions than their counterparts in NoAI-NoAI. The treatment thus affects not only the level of unconditional cooperation but also the strength of conditional

reciprocity among those who exhibit it.

Turning to the WLG, we classify participants as conditional coordinators using the same approach as in the PGG. The share of conditional coordinators is similar across treatment conditions and is higher than the share of conditional cooperators in the PGG, which is consistent with the coordination incentives embedded in the WLG. The conditional coordination patterns are almost identical across treatment conditions, as shown in the bottom row of Figure 2. Accordingly, we find no significant differences in either slopes or levels across conditions (see Column 6 of Table OA-2). Comparing behavioral types across games, 68.3% of participants can be classified as both conditional cooperators in the PGG and conditional coordinators in the WLG, with only minor differences across treatment conditions.

Table 2 also provides evidence on how the possibility that other group members may use AI affects participants' normative and empirical beliefs. In the PGG, first-order normative beliefs—what the participant finds to be a socially appropriate contribution to the public good—are lower in NoAI-AI than in NoAI-NoAI. Similarly, they also exhibit a smaller descriptive belief, i.e., they expect others to contribute, on average, slightly less. However, neither difference is statistically significant. The clearest difference emerges for second-order normative beliefs: participants in NoAI-AI report significantly lower beliefs about what others consider an appropriate contribution than participants in NoAI-NoAI (8.64 compared to 6.61). This suggests that the possibility that others may use AI primarily affects the social norm of appropriate behavior in the PGG.

In the WLG, by contrast, treatment differences are more pronounced for first-order normative and descriptive beliefs. Participants in NoAI-AI report significantly lower first-order normative beliefs than participants in NoAI-NoAI (10.98 compared to 8.85), and they also expect lower efforts by other group members (11.38 compared to 9.41). Second-order normative beliefs are lower in NoAI-AI as well, but the difference is not statistically significant.

Overall, these results suggest that the possibility that others may use AI lowers beliefs about both appropriate and expected contribution or effort decisions, although the relevant belief dimension differs between the cooperation and coordination settings. These shifts in beliefs are consistent with the observed treatment effects on unconditional behavior: when participants expect others to contribute less or exert less effort, or perceive the relevant social norms as less demanding, they may themselves become less willing to cooperate or coordinate unconditionally. The belief channel may therefore partly mediate the negative indirect effect of others' possible AI access on unconditional contributions and effort provision.

4 Discussion and conclusion

In this paper, we studied whether the possibility that others may use AI affects human decision-making even when the decision-maker herself is not allowed to use AI. Participants in our experiment played a one-shot PGG and a one-shot WLG. The experimental design allowed us to causally identify the indirect effect of others' possible AI access on cooperation and coordination behavior.

We found that others' possible AI access had negative indirect effects on the cooperation and coordination behavior of participants who were not allowed to use AI themselves. In the PGG, these participants contributed less when their group members could use AI; in

the WLG, they chose lower effort levels. Thus, rather than observing opposite effects in cooperation and coordination settings, as hypothesized *ex ante*, we found a common negative effect of others' possible AI access on both unconditional cooperation and effort provision. This extends existing evidence that AI involvement affects behavior when AI directly replaces a human decision-maker (Dell'Acqua et al., 2025; Dvorak et al., 2025). It seems that the effect does not require direct replacement. The mere possibility that other group members may consult AI seems sufficient to affect behavior, beliefs, and perceived closeness—which may be particularly relevant in real-world settings, where AI use by others is rarely observable or verifiable. At the same time, we found that conditional decision profiles remained largely stable: the shares of conditional cooperators and conditional coordinators did not differ significantly across treatments, and conditional coordination patterns in the weakest-link game were almost identical across treatments. This suggests that participants' fundamental cooperation and coordination types were largely unaffected by our treatment—which underlines previous evidence from situations that do not involve AI that cooperation types seem to be relatively stable individual characteristics persisting across contexts and populations (Thöni & Volk, 2018).

Why do the effects in the PGG and the WLG point in the same direction? One interpretation is that participants may not have perceived the weakest-link game purely as a coordination problem. Formally, the two games differ: the PGG involves a conflict between individual and collective incentives, whereas the WLG involves multiple equilibria and strategic uncertainty about the minimum effort. Nevertheless, the WLG also has a strong team-production and mutual-dependence component. A participant's costly effort benefits the group only if others also exert sufficient effort; otherwise, the participant bears costs without generating the desired group payoff. In this sense, the weakest-link game may be experienced psychologically as a cooperation problem, even if its formal game-theoretic structure is one of coordination. This interpretation is consistent with previous work that relates weakest-link mechanisms to reciprocity, matching, and conditional cooperation (Croson et al., 2005), as well as work comparing punishment in cooperation and coordination games (Albrecht & Kube, 2018). Under this perspective, knowing that some group members may use AI could reduce effort provision for reasons similar to those that reduce contributions in the public goods game: participants may perceive AI-assisted group members as less socially close, less normatively committed, or more likely to act in a detached and payoff-oriented manner.

Our measure of subjective closeness supports this interpretation. Participants who were not allowed to use AI reported lower closeness to group members who might use AI than to group members who could not use AI. This suggests that AI access does not only change beliefs about strategic sophistication; it may also change the perceived social relationship among group members. Such social distance can matter in both games. In the public goods game, lower perceived closeness may reduce social obligation and reciprocal cooperation. In the weakest-link game, it may reduce confidence that others will choose the high-effort path needed for successful team production. These interpretations are consistent with evidence that group cohesion is strongly associated with performance in weak-link team production (Gächter et al., 2025), that reduced social distance raises contributions in public goods settings (Eckel & Grossman, 2005), and that social distance shapes cooperative behavior in iterated dilemma games (Zhang et al., 2022a).

A second explanation is that others' possible AI access may have increased rather than decreased strategic uncertainty. Our preregistered hypothesis was based on the idea that AI assistance could help group members identify the payoff-dominant equilibrium. However, participants did not know which AI tools others might use, whether they would use them at all, how they would prompt them, or what advice they would receive—and our auxiliary simulations illustrate that such uncertainty is plausible: ChatGPT 5.3 Instant and Gemini 3 Fast produced sharply different recommendations in both games. This heterogeneity is consistent with broader evidence that different LLMs exhibit systematically divergent strategic behavior even in identical game-theoretic settings (Lorè & Heydari, 2024; Xie et al., 2024). Others' possible AI access therefore differs from the coordination-enhancing mechanisms studied in prior work, which create common expectations through shared signals, communication, or institutional changes (Brandts & Cooper, 2006, 2007; Blume & Ortmann, 2007; Kogan et al., 2011). Instead, it changes what participants know about the decision-making process of others, introducing a novel source of strategic uncertainty. In the weakest-link game, where the payoff from high effort depends on the lowest effort chosen by any group member, even small doubts about others' choices can make lower effort attractive (Huyck et al., 1990; Devetag & Ortmann, 2007; Hanaki et al., 2016).

Several limitations should be noted. First, our main treatment varies others' possible AI access, not verified AI use. This is intentional, as many real-world settings involve uncertainty about whether others use AI. However, the treatment effect should therefore be interpreted as the effect of anticipated or possible AI assistance by others, not the effect of actual AI-generated advice. Second, we study one-shot interactions. Repeated interaction may change the effect of AI access as participants learn whether AI-assisted group members are reliable, selfish, cooperative, or predictable. Third, our design does not identify which specific belief channel is causal. We observe changes in perceived closeness, normative beliefs, and descriptive beliefs, but these measures are post-treatment outcomes. Future work could experimentally manipulate perceived AI advice, perceived social distance, or information about the content of AI recommendations to disentangle these mechanisms more directly. Finally, AI systems evolve quickly. The advice generated by current models, and beliefs about such advice, may differ from those generated by future systems.

Moreover, future research could extend the present design in several directions. One natural extension is to compare uncertain AI access with disclosed actual AI use and with disclosed AI-generated advice. This would help separate the effects of uncertainty, source identity, and advice content. A second extension is to study repeated games, where participants can learn about the behavior of AI-assisted group members over time. A third extension is to investigate institutional remedies. In public goods games, punishment, rewards, communication, and information provision can sustain cooperation; in weakest-link games, communication, focal recommendations, and shared histories can facilitate efficient coordination. It remains open whether these mechanisms work similarly when some group members may use AI. For example, transparent disclosure rules, common AI recommendations, shared prompts, or symmetric access to AI tools might reduce uncertainty and restore coordination, whereas asymmetric and uncertain access to AI advice might continue to undermine trust and perceived closeness.

Overall, our findings show that AI can affect strategic interaction even when it is not directly used by the decision-maker. Merely knowing that others may consult AI reduces uncon-

ditional cooperation in a public goods game and unconditional effort provision in a weakest-link game. These effects are accompanied by lower perceived social closeness and shifts in beliefs about appropriate and expected behavior, while conditional cooperation and coordination profiles remain largely stable. This suggests that AI is not only a private decision aid; it can also become part of the social and strategic environment. As AI tools become increasingly available, their behavioral consequences may therefore extend beyond users themselves to those who interact with, compete with, or depend on people who may use them.

References

- Akata, E., Schulz, L., Coda-Forno, J., Oh, S. J., Bethge, M., & Schulz, E. (2025). Playing repeated games with large language models. *Nature Human Behaviour*, 9(7), 1380–1390. <https://doi.org/10.1038/s41562-025-02172-y>.
- Albrecht, F., & Kube, S. (2018). Peer-Punishment in a Cooperation and a Coordination Game. *Games*, 9(3), 54. <https://doi.org/10.3390/g9030054>.
- Angelovski, A., Di Cagno, D., Güth, W., Marazzi, F., & Panaccione, L. (2018). Behavioral spillovers in local public good provision: An experimental study. *Journal of Economic Psychology*, 67, 116–134. <https://doi.org/10.1016/j.joep.2018.05.003>.
- Baader, M., Starmer, C., Tufano, F., & Gächter, S. (2024). Introducing IOS11 as an extended interactive version of the ‘Inclusion of Other in the Self’ scale to estimate relationship closeness. *Scientific Reports*, 14(1), 8901. <https://doi.org/10.1038/s41598-024-58042-6>.
- Barak, D., & Costa-Gomes, M. (2025). *Humans expect rationality and cooperation from LLM opponents in strategic games* (arXiv Preprint). <https://arxiv.org/abs/2505.11011>.
- Ben-Michael, E., Greiner, D. J., Huang, M., Imai, K., Jiang, Z., & Shin, S. (2025). Does AI help humans make better decisions? A statistical evaluation framework for experimental and observational studies. *Proceedings of the National Academy of Sciences*, 122(38), e2505106122. <https://doi.org/10.1073/pnas.2505106122>.
- Bicchieri, C. (2006). *The Grammar of Society: The nature and dynamics of social norms*. Cambridge University Press.
- Blume, A., & Ortmann, A. (2007). The effects of costless pre-play communication: Experimental evidence from games with Pareto-ranked equilibria. *Journal of Economic Theory*, 132(1), 274–290. <https://doi.org/10.1016/j.jet.2005.03.001>.
- Bock, O., Baetge, I., & Nicklisch, A. (2014). Hroot: Hamburg Registration and Organization Online Tool. *European Economic Review*, 71, 117–120. <https://doi.org/10.1016/j.eurocorev.2014.07.003>.
- Bortolotti, S., Devetag, G., & Ortmann, A. (2016). Group incentives or individual incentives? A real-effort weak-link experiment. *Journal of Economic Psychology*, 56, 60–73. <https://doi.org/10.1016/j.joep.2016.05.004>.
- Brandts, J., & Cooper, D. J. (2006). A Change Would Do You Good An Experimental Study on How to Overcome Coordination Failure in Organizations. *American Economic Review*, 96(3), 669–693. <https://doi.org/10.1257/aer.96.3.669>.
- Brandts, J., & Cooper, D. J. (2007). It’s What You Say, Not What You Pay: An Experimental Study of Manager-Employee Relationships in Overcoming Coordination Failure. *Journal of*

- the European Economic Association*, 5(6), 1223–1268. <https://doi.org/10.1162/JEEA.2007.5.6.1223>.
- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at Work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>.
- Capraro, V., Di Paolo, R., & Pizziol, V. (2025). A publicly available benchmark for assessing large language models' ability to predict how humans balance self-interest and the interest of others. *Scientific Reports*, 15(1), 21428. <https://doi.org/10.1038/s41598-025-01715-7>.
- Catola, M., D'Alessandro, S., Guarnieri, P., & Pizziol, V. (2023). Multilevel public goods game: Levelling up, substitution and crowding-in effects. *Journal of Economic Psychology*, 97, 102626. <https://doi.org/10.1016/j.joep.2023.102626>.
- Chaudhuri, A. (2011). Sustaining cooperation in laboratory public goods experiments: A selective survey of the literature. *Experimental Economics*, 14(1), 47–83. <https://doi.org/10.1007/s10683-010-9257-1>.
- Chaudhuri, A., Paichayontvijit, T., & Smith, A. (2017). Belief heterogeneity and contributions decay among conditional cooperators in public goods games. *Journal of Economic Psychology*, 58, 15–30. <https://doi.org/10.1016/j.joep.2016.11.004>.
- Chaudhuri, A., Paichayontvijit, T., & So, T. (2015). Team versus individual behavior in the minimum effort coordination game. *Journal of Economic Psychology*, 47, 85–102. <https://doi.org/10.1016/j.joep.2015.02.002>.
- Choi, J.-K., & Ahn, T. (2013). Strategic reward and altruistic punishment support cooperation in a public goods game experiment. *Journal of Economic Psychology*, 35, 17–30. <https://doi.org/10.1016/j.joep.2013.01.001>.
- Chugunova, M., & Luhan, W. J. (2025). Ruled by robots: Preference for algorithmic decision makers and perceptions of their choices. *Public Choice*, 202(1-2), 1–24. <https://doi.org/10.1007/s11127-024-01178-w>.
- Chugunova, M., & Sele, D. (2022). We and It: An interdisciplinary review of the experimental evidence on how humans interact with machines. *Journal of Behavioral and Experimental Economics*, 99, 101897. <https://doi.org/10.1016/j.socec.2022.101897>.
- Crandall, J. W., Oudah, M., Tennom, Ishowo-Oloko, F., Abdallah, S., Bonnefon, J.-F., Cebrian, M., Shariff, A., Goodrich, M. A., & Rahwan, I. (2018). Cooperating with machines. *Nature Communications*, 9(1), 233. <https://doi.org/10.1038/s41467-017-02597-8>.
- Croson, R., Fatas, E., & Neugebauer, T. (2005). Reciprocity, matching and conditional cooperation in two public goods games. *Economics Letters*, 87(1), 95–101. <https://doi.org/10.1016/j.econlet.2004.10.007>.
- Dell'Acqua, E., Kogut, B., & Perkowski, P. (2025). Super Mario Meets AI: Experimental Effects of Automation and Skills on Team Performance and Coordination. *Review of Economics and Statistics*, 107(4), 951–966. https://doi.org/10.1162/rest_a_01328.
- Devetag, G., & Ortmann, A. (2007). When and why? A critical survey on coordination failure in the laboratory. *Experimental Economics*, 10(3), 331–344. <https://doi.org/10.1007/s10683-007-9178-9>.
- Drouvelis, M., Nosenzo, D., & Sefton, M. (2017). Team incentives and leadership. *Journal of Economic Psychology*, 62, 173–185. <https://doi.org/10.1016/j.joep.2017.07.002>.

- Dvorak, F., Stumpf, R., Fehrler, S., & Fischbacher, U. (2025). Adverse reactions to the use of large language models in social interactions. *PNAS Nexus*, 4(4), pgaf112. <https://doi.org/10.1093/pnasnexus/pgaf112>.
- Eckel, C. C., & Grossman, P. J. (2005). Managing diversity by creating team identity. *Journal of Economic Behavior & Organization*, 58(3), 371–392. <https://doi.org/10.1016/j.jebo.2004.01.003>.
- Engel, C., Kube, S., & Kurschilgen, M. (2021). Managing expectations: How selective information affects cooperation and punishment in social dilemma games. *Journal of Economic Behavior & Organization*, 187, 111–136. <https://doi.org/10.1016/j.jebo.2021.04.029>.
- Engl, F., Riedl, A., & Weber, R. (2021). Spillover Effects of Institutions on Cooperative Behavior, Preferences, and Beliefs. *American Economic Journal: Microeconomics*, 13(4), 261–299. <https://doi.org/10.1257/mic.20180336>.
- Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., & Sunde, U. (2018). Global Evidence on Economic Preferences. *The Quarterly Journal of Economics*, 133(4), 1645–1692. <https://doi.org/10.1093/qje/qjy013>.
- Fischbacher, U., & Gächter, S. (2010). Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Goods Experiments. *American Economic Review*, 100(1), 541–556. <https://doi.org/10.1257/aer.100.1.541>.
- Fischbacher, U., Gächter, S., & Fehr, E. (2001). Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters*, 71(3), 397–404. [https://doi.org/10.1016/S0165-1765\(01\)00394-9](https://doi.org/10.1016/S0165-1765(01)00394-9).
- Fischbacher, U., Gächter, S., & Quercia, S. (2012). The behavioral validity of the strategy method in public good experiments. *Journal of Economic Psychology*, 33(4), 897–913. <https://doi.org/10.1016/j.joep.2012.04.002>.
- Gächter, S., Kölle, F., & Quercia, S. (2017). Reciprocity and the tragedies of maintaining and providing the commons. *Nature Human Behaviour*, 1(9), 650–656. <https://doi.org/10.1038/s41562-017-0191-5>.
- Gächter, S., Starmer, C., Thöni, C., Tufano, F., & Weber, T. O. (2022). Social closeness can help, harm and be irrelevant in solving pure coordination problems. *Economics Letters*, 216, 110552. <https://doi.org/10.1016/j.econlet.2022.110552>.
- Gächter, S., Starmer, C., & Tufano, F. (2025). Measuring Group Cohesion to Reveal the Power of Social Relationships in Team Production. *Review of Economics and Statistics*, 107(2), 539–554. https://doi.org/10.1162/rest_a_01283.
- Gächter, S., & Thöni, C. (2005). Social Learning and Voluntary Cooperation Among Like-Minded People. *Journal of the European Economic Association*, 3(2/3), 303–314. <http://www.jstor.org/stable/40004973>.
- Gieselmann, M., Erdsiek, D., Rost, V., & Sassenberg, K. (2025). Do managers accept artificial intelligence? Insights into the role of business area and AI functionality. *Journal of Economic Psychology*, 108, 102804. <https://doi.org/10.1016/j.joep.2025.102804>.
- Goeschl, T., & Jarke, J. (2016). Second and third party punishment under costly monitoring. *Journal of Economic Psychology*, 54, 124–133. <https://doi.org/10.1016/j.joep.2016.03.002>.

- Greiner, B., Grünwald, P., Lindner, T., Lintner, G., & Wiernsperger, M. (2025). Incentives, Framing, and Reliance on Algorithmic Advice: An Experimental Study. *Management Science*, 72(1), 302–322. <https://doi.org/10.1287/mnsc.2022.02777>.
- Grieco, D., Faillo, M., & Zarri, L. (2017). Enforcing cooperation in public goods games: Is one punisher enough? *Journal of Economic Psychology*, 61, 55–73. <https://doi.org/10.1016/j.joep.2017.03.007>.
- Guala, F., Mittone, L., & Ploner, M. (2013). Group membership, team preferences, and expectations. *Journal of Economic Behavior & Organization*, 86, 183–190. <https://doi.org/10.1016/j.jebo.2012.12.003>.
- Guillen, P., Fatas, E., & Brañas-Garza, P. (2010). Inducing efficient conditional cooperation patterns in public goods games, an experimental investigation. *Journal of Economic Psychology*, 31(6), 872–883. <https://doi.org/10.1016/j.joep.2010.07.002>.
- Hanaki, N., Jacquemet, N., Luchini, S., & Zylbersztejn, A. (2016). Cognitive ability and the effect of strategic uncertainty. *Theory and Decision*, 81(1), 101–121. <https://doi.org/10.1007/s11238-015-9525-9>.
- Huseynov, S. (2025). ChatGPT and the labor market: Unraveling the effect of AI discussions on students' earning expectations. *Journal of Economic Psychology*, 108, 102803. <https://doi.org/10.1016/j.joep.2025.102803>.
- Huyck, J. B. V., Battalio, R. C., & Beil, R. O. (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *The American Economic Review*, 80(1), 234–248. <http://www.jstor.org/stable/2006745>.
- Ishowo-Oloko, F., Bonnefon, J.-F., Soroye, Z., Crandall, J., Rahwan, I., & Rahwan, T. (2019). Behavioural evidence for a transparency–efficiency tradeoff in human–machine cooperation. *Nature Machine Intelligence*, 1(11), 517–521. <https://doi.org/10.1038/s42256-019-0113-5>.
- Jiang, G., Yang, S., Wang, Y., & Hui, P. (2026). When trust collides: Exploring human-LLM cooperation intention through the prisoner's dilemma. *International Journal of Human-Computer Studies*, 209, 103740. <https://doi.org/10.1016/j.ijhcs.2026.103740>.
- Jobu Babin, J., & Chauhan, H. S. (2026). Chatbot or humanaut? How the source of advice impacts prosocial behavior. *Journal of Behavioral and Experimental Economics*, 121, 102509. <https://doi.org/10.1016/j.socec.2026.102509>.
- Jouxte, J. (2019). Voluntary contributions of time: Time-based incentives in a linear public goods game. *Journal of Economic Psychology*, 75, 102139. <https://doi.org/10.1016/j.joep.2019.01.002>.
- Kaplan, J. (2016). *Artificial intelligence: What everyone needs to know*. Oxford University Press.
- Karpus, J., Krüger, A., Verba, J. T., Bahrami, B., & Deroy, O. (2021). Algorithm exploitation: Humans are keen to exploit benevolent AI. *iScience*, 24(6), 102679. <https://doi.org/10.1016/j.isci.2021.102679>.
- Karpus, J., Shirai, R., Verba, J. T., Schulte, R., Weigert, M., Bahrami, B., Watanabe, K., & Deroy, O. (2025). Human cooperation with artificial agents varies across countries. *Scientific Reports*, 15(1), 10000. <https://doi.org/10.1038/s41598-025-92977-8>.
- Köbis, N., Bonnefon, J.-F., & Rahwan, I. (2021). Bad machines corrupt good morals. *Nature Human Behaviour*, 5(6), 679–685. <https://doi.org/10.1038/s41562-021-01128-2>.

- Kocher, M. G., Martinsson, P., Matzat, D., & Wollbrant, C. (2015). The role of beliefs, trust, and risk in contributions to a public good. *Journal of Economic Psychology*, *51*, 236–244. <https://doi.org/10.1016/j.joep.2015.10.001>.
- Kogan, S., Kwasnica, A. M., & Weber, R. A. (2011). Coordination in the Presence of Asset Markets. *American Economic Review*, *101*(2), 927–947. <https://doi.org/10.1257/aer.101.2.927>.
- Kölle, F., & Quercia, S. (2021). The influence of empirical and normative expectations on cooperation. *Journal of Economic Behavior & Organization*, *190*, 691–703. <https://doi.org/10.1016/j.jebo.2021.08.018>.
- Krupka, E. L., & Weber, R. A. (2013). Identifying Social Norms Using Coordination Games: Why Does Dictator Game Sharing Vary? *Journal of the European Economic Association*, *11*(3), 495–524. <https://doi.org/10.1111/jeea.12006>.
- Leib, M., Köbis, N., Rilke, R. M., Hagens, M., & Irlenbusch, B. (2024). Corrupted by Algorithms? How AI-generated and Human-written Advice Shape (Dis)honesty. *The Economic Journal*, *134*(658), 766–784. <https://doi.org/10.1093/ej/uead056>.
- Leib, M., Köbis, N., & Soraperra, I. (2025). Does AI and human advice mitigate punishment for selfish behavior? An experiment on AI ethics from a psychological perspective. *Computers in Human Behavior*, *171*, 108709. <https://doi.org/10.1016/j.chb.2025.108709>.
- Lorè, N., & Heydari, B. (2024). Strategic behavior of large language models and the role of game structure versus contextual framing. *Scientific Reports*, *14*(1), 18490. <https://doi.org/10.1038/s41598-024-69032-z>.
- March, C. (2021). Strategic interactions between humans and artificial intelligence: Lessons from experiments with computer players. *Journal of Economic Psychology*, *87*, 102426. <https://doi.org/10.1016/j.joep.2021.102426>.
- Mei, Q., Xie, Y., Yuan, W., & Jackson, M. O. (2024). A Turing test of whether AI chatbots are behaviorally similar to humans. *Proceedings of the National Academy of Sciences*, *121*(9), e2313925121. <https://doi.org/10.1073/pnas.2313925121>.
- Morgan, S. N., Mason, N. M., & Shupp, R. S. (2019). The effects of voice with(out) punishment: Public goods provision and rule compliance. *Journal of Economic Psychology*, *74*, 102190. <https://doi.org/10.1016/j.joep.2019.102190>.
- Myers, S., & Everett, J. A. (2025). People expect artificial moral advisors to be more utilitarian and distrust utilitarian moral advisors. *Cognition*, *256*, 106028. <https://doi.org/10.1016/j.cognition.2024.106028>.
- Neugebauer, T., Perote, J., Schmidt, U., & Loos, M. (2009). Selfish-biased conditional cooperation: On the decline of contributions in repeated public goods experiments. *Journal of Economic Psychology*, *30*(1), 52–60. <https://doi.org/10.1016/j.joep.2008.04.005>.
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, *381*(6654), 187–192. <https://doi.org/10.1126/science.adh2586>.
- Orchard, T., & Tasiemski, L. (2023). The rise of generative AI and possible effects on the economy. *Economics and Business Review*, *9*(2), 9–26. <https://doi.org/10.18559/eb.2023.2.732>.

- R Core Team. (2022). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>.
- Rustagi, D., Engel, S., & Kosfeld, M. (2010). Conditional Cooperation and Costly Monitoring Explain Success in Forest Commons Management. *Science*, 330(6006), 961–965. <https://doi.org/10.1126/science.1193649>.
- Selten, R. (1967). Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperimentes. In H. Sauermann (Ed.), *Beiträge zur experimentellen wirtschaftsforschung* (pp. 136–168, Vol. 1). Mohr.
- Selten, R. (1998). Axiomatic characterization of the quadratic scoring rule. *Experimental Economics*, 1, 43–61. <https://doi.org/10.1023/A:1009957816843>.
- Shen, C., Karpus, J., Kosch, T., Fernandes, D., Mello, B., Welsch, R., & Villa, S. (2025). The Impact of Asymmetric AI Assistance on Decision-Making in Social Dilemmas: A Study on Human Augmentation in Economic Games. *Proceedings of the Augmented Humans International Conference 2025*, 187–198. <https://doi.org/10.1145/3745900.3746084>.
- Shirado, H., & Christakis, N. A. (2017). Locally noisy autonomous agents improve global human coordination in network experiments. *Nature*, 545(7654), 370–374. <https://doi.org/10.1038/nature22332>.
- Stein, J.-P., Messingschlager, T., Gnambs, T., Hutmacher, F., & Appel, M. (2024). Attitudes towards AI: Measurement and associations with personality. *Scientific Reports*, 14, 2909. <https://doi.org/10.1038/s41598-024-53335-2>.
- Suri, D., Gächter, S., & Kube, S. (2026a). *AI versus humans as authority figures: Evidence from a rule-compliance experiment* (CeDEx Discussion Paper No. 2026-02).
- Suri, D., Gächter, S., Kube, S., & Schultz, J. (2026b). *The resilience of rule compliance in a polarized society* (CeDEx Discussion Paper No. 2026-01).
- Tan, F., & Xiao, E. (2018). Third-party punishment: Retribution or deterrence? *Journal of Economic Psychology*, 67, 34–46. <https://doi.org/10.1016/j.joep.2018.03.003>.
- Teutloff, O., Einsiedler, J., Kässi, O., Braesemann, F., Mishkin, P., & Del Rio-Chanona, R. M. (2025). Winners and losers of generative AI: Early Evidence of Shifts in Freelancer Demand. *Journal of Economic Behavior & Organization*, 235, 106845. <https://doi.org/10.1016/j.jebo.2024.106845>.
- Thöni, C., & Volk, S. (2018). Conditional cooperation: Review and refinement. *Economics Letters*, 171, 37–40. <https://doi.org/10.1016/j.econlet.2018.06.022>.
- Tse, T. T. K., Hanaki, N., & Mao, B. (2024). Beware the performance of an algorithm before relying on it: Evidence from a stock price forecasting experiment. *Journal of Economic Psychology*, 102, 102727. <https://doi.org/10.1016/j.joep.2024.102727>.
- Volk, S., Thöni, C., & Ruigrok, W. (2012). Temporal stability and psychological foundations of cooperation preferences. *Journal of Economic Behavior & Organization*, 81(2), 664–676. <https://doi.org/10.1016/j.jebo.2011.10.006>.
- Weber, R. A. (2006). Managing Growth to Achieve Efficient Coordination in Large Groups. *American Economic Review*, 96(1), 114–126. <https://doi.org/10.1257/000282806776157588>.
- Wester, J., De Jong, S., Pohl, H., & Van Berkel, N. (2024). Exploring people's perceptions of LLM-generated advice. *Computers in Human Behavior: Artificial Humans*, 2(2), 100072. <https://doi.org/10.1016/j.chbah.2024.100072>.

- Xie, Y., Liu, Y., Ma, Z., Shi, L., Wang, X., Yuan, W., Jackson, M. O., & Mei, Q. (2024). *How Different AI Chatbots Behave? Benchmarking Large Language Models in Behavioral Economics Games* (arXiv Preprint). <https://arxiv.org/abs/2412.12362>.
- Xie, Y., Mei, Q., Yuan, W., & Jackson, M. O. (2025). Using large language models to categorize strategic situations and decipher motivations behind human behaviors. *Proceedings of the National Academy of Sciences*, *122*(35), e2512075122. <https://doi.org/10.1073/pnas.2512075122>.
- Yao, J., Wang, K., Hsieh, R., Zhou, H., Zou, T., Cheng, Z., Wang, Z., & Viswanath, P. (2025). *SPIN-Bench: How Well Do LLMs Plan Strategically and Reason Socially?* (arXiv Preprint). <https://arxiv.org/abs/2503.12349>.
- Yin, J., Ngiam, K. Y., Tan, S. S.-L., & Teo, H. H. (2025). Designing AI-Based Work Processes: How the Timing of AI Advice Affects Diagnostic Decision Making. *Management Science*, *71*(11), 9361–9383. <https://doi.org/10.1287/mnsc.2022.01454>.
- Zelmer, J. (2003). Linear Public Goods Experiments: A Meta-Analysis. *Experimental Economics*, *6*(3), 299–310. <https://doi.org/10.1023/A:1026277420119>.
- Zhang, H. (2019). Common fate motivates cooperation: The influence of risks on contributions to public goods. *Journal of Economic Psychology*, *70*, 12–21. <https://doi.org/10.1016/j.joep.2018.10.012>.
- Zhang, L., Jin, Y., Xia, L., Xu, B., & Syed Abdullah, S. M. (2022a). The Effects of Social Distance and Asymmetric Reward and Punishment on Individual Cooperative Behavior in Dilemma Situations. *Frontiers in Psychology*, *13*, 816168. <https://doi.org/10.3389/fpsyg.2022.816168>.
- Zhang, Z., Chen, Z., & Xu, L. (2022b). Artificial intelligence and moral dilemmas: Perception of ethical decision-making in AI. *Journal of Experimental Social Psychology*, *101*, 104327. <https://doi.org/10.1016/j.jesp.2022.104327>.

Online Appendix
for
Cooperation and coordination when others may use AI

Dominik Atella-Suri*, Sebastian Kube

*Corresponding author. E-mail: dsuri@uni-bonn.de.

Contents

OA-1 Supplementary Figures	2
Figure OA-1: Subjective closeness to group members in AI-AI	2
Figure OA-2: Conditional contribution types in AI-AI	2
Figure OA-3: Raw data of the public goods game by treatment condition	3
Figure OA-4: Raw data of the weakest-link game by treatment condition	4
OA-2 Supplementary Tables	5
Table OA-1: Treatment effects in the public goods game	5
Table OA-2: Treatment effects in the weakest-link game	6
OA-3 Experimental Instructions	7

OA-1 Supplementary Figures




Figure OA-1: Subjective closeness to group members in AI-AI. The figure shows the distribution of subjective closeness toward other group members that also may use AI to inform their decisions in treatment condition AI-AI. Subjective closeness is measured by the IOS11 score, where 1 means subjectively very distant and 11 means subjectively very close. The dashed lines represent the respective means.

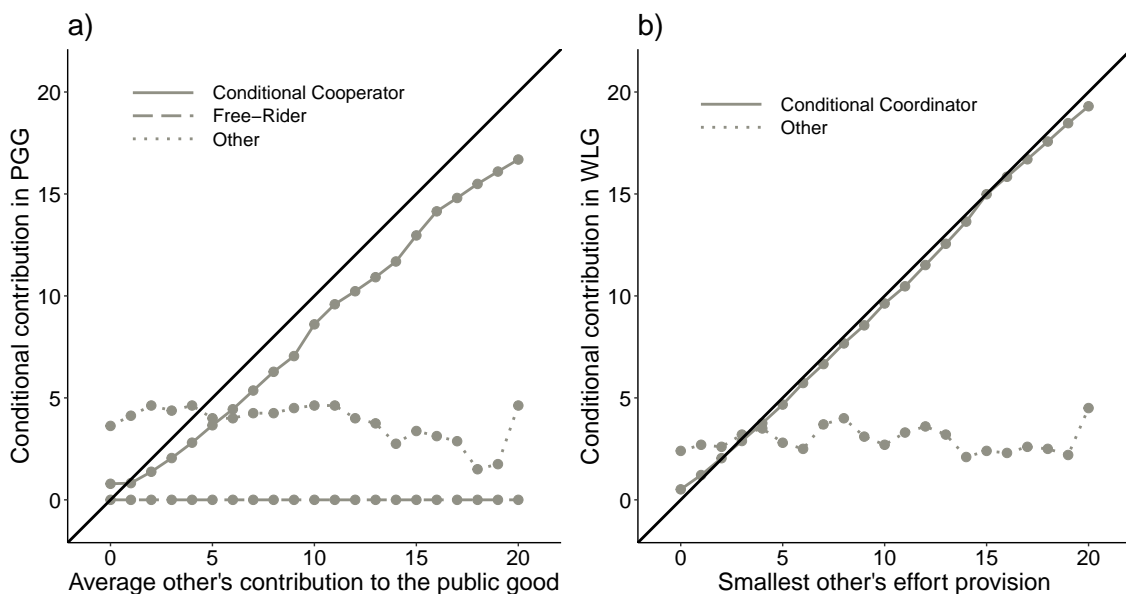


Figure OA-2: Conditional contribution types in AI-AI. The figure shows the conditional contribution types in treatment condition AI-AI. The solid lines represent the 45 degree line.

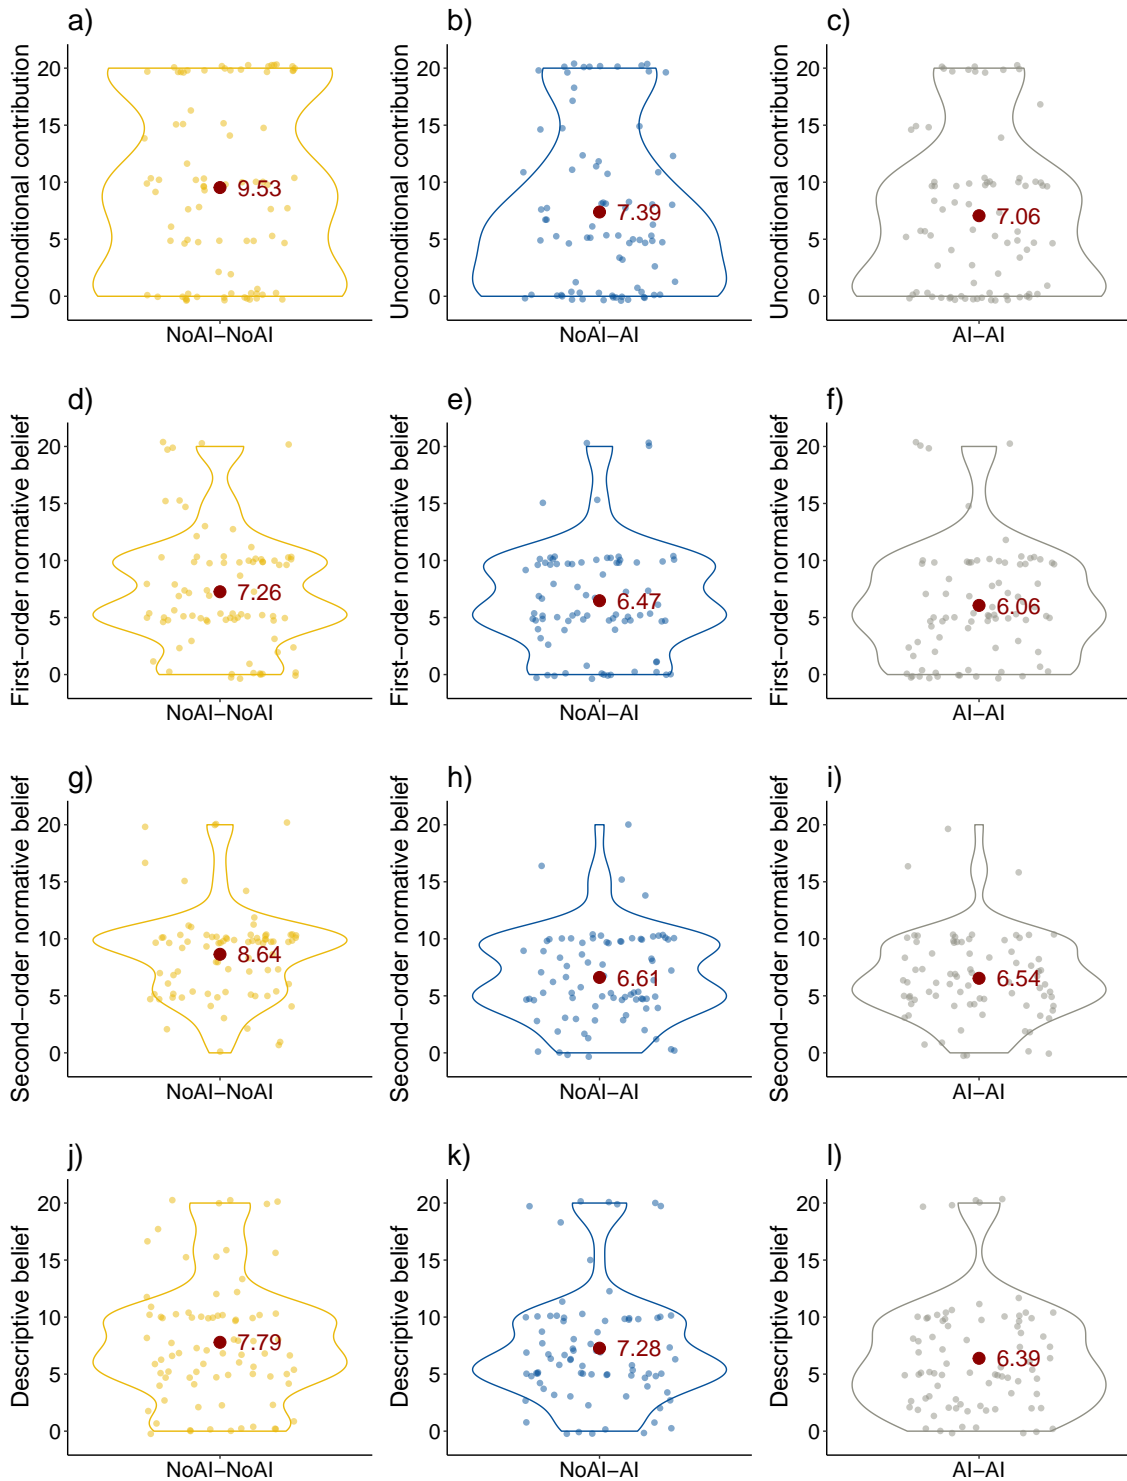


Figure OA-3: Raw data of the public goods game by treatment condition. The figure shows the raw data of the public goods game separately for each treatment condition in violin plots. The first row shows the unconditional contribution decision in NoAI-NoAI (left), NoAI-AI (middle), and AI-AI (right). The second row from the top shows answers to the first-order normative belief question. The third row from the top shows answers to the second-order normative belief question. The bottom row shows answers to the descriptive belief question. The red dots indicate mean values.

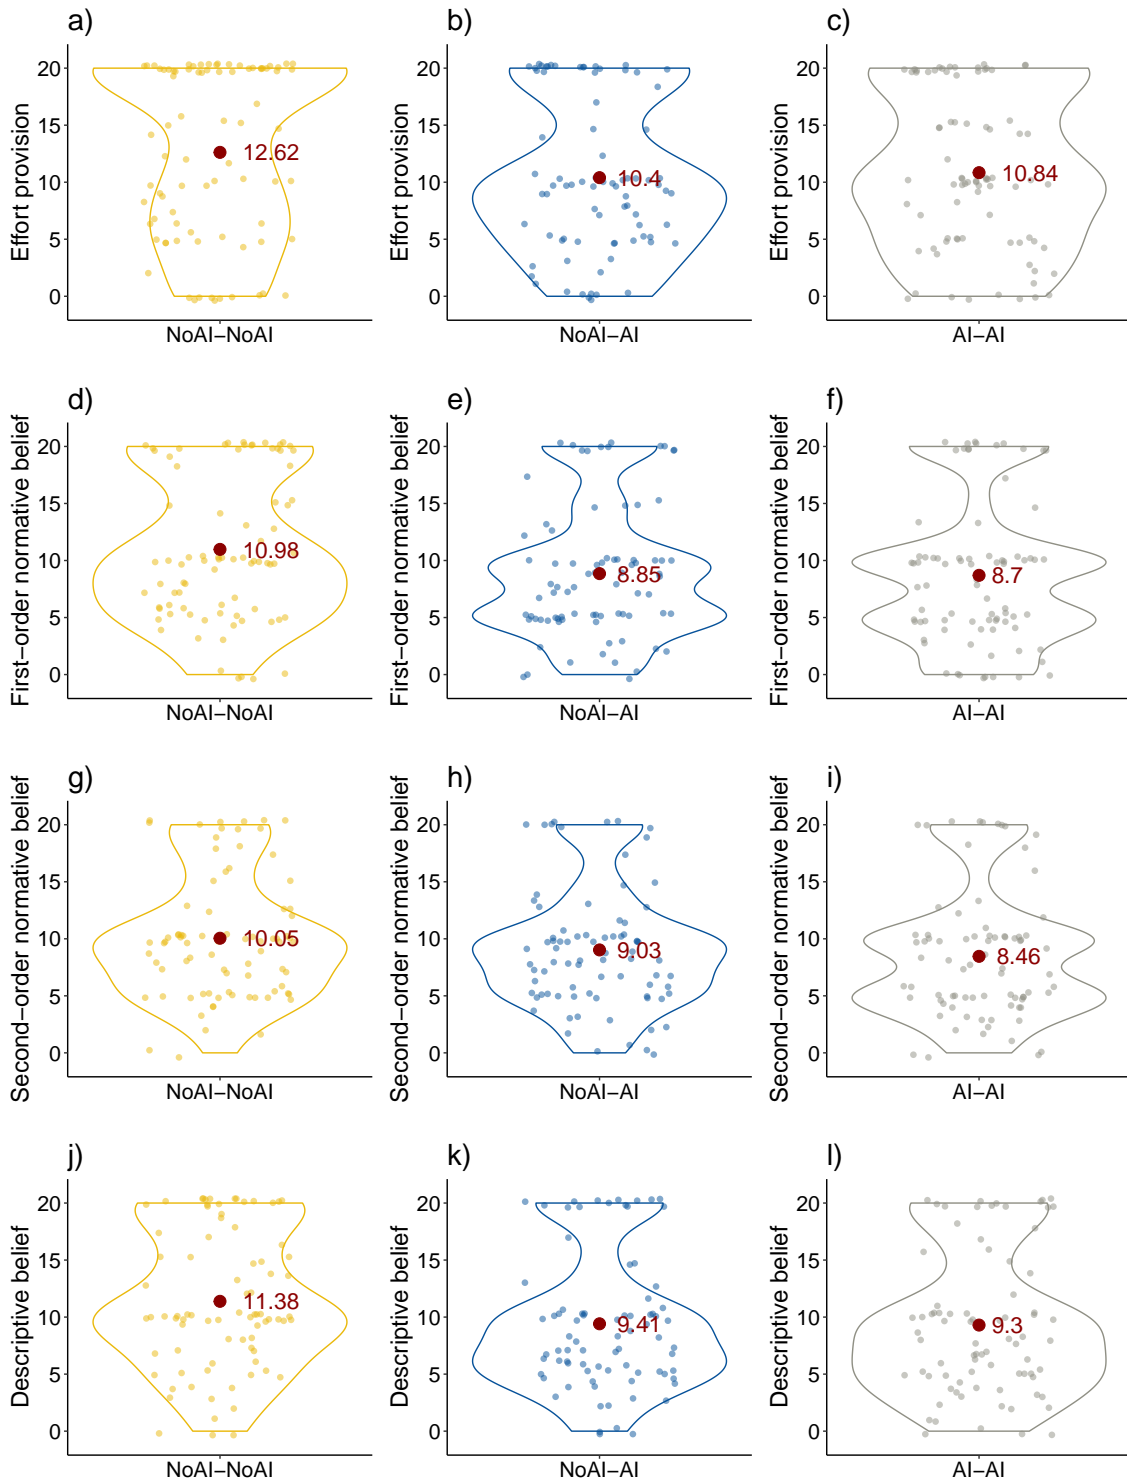


Figure OA-4: Raw data of the weakest-link game by treatment condition. The figure shows the raw data of the weakest-link game separately for each treatment condition in violin plots. The first row shows the effort provision decision in NoAI-NoAI (left), NoAI-AI (middle), and AI-AI (right). The second row from the top shows answers to the first-order normative belief question. The third row from the top shows answers to the second-order normative belief question. The bottom row shows answers to the descriptive belief question. The red dots indicate mean values.

OA-2 Supplementary Tables

Table OA-1: Treatment effects in the public goods game

	<i>Dependent variable:</i>					
	<i>Unconditional contribution</i>			<i>Conditional contribution</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	3.93** (1.20)	-1.66 (3.27)	-2.65 (3.72)	-7.02** (2.49)	-1.64 (2.18)	-4.92* (2.30)
NoAI-AI	-2.53** (0.86)	-2.23** (0.82)	-2.21** (0.82)	0.28 (0.68)	0.93* (0.40)	0.75* (0.38)
1st-order norm. belief	0.74*** (0.11)	0.72*** (0.12)	0.73*** (0.12)			
2nd-order norm. belief	-0.67*** (0.15)	-0.71*** (0.14)	-0.70*** (0.15)			
Descriptive belief	0.78*** (0.10)	0.74*** (0.10)	0.72*** (0.10)			
Others' Contribution				0.69*** (0.05)	0.94*** (0.03)	0.94*** (0.03)
IOS11		0.35* (0.14)	0.35* (0.14)	0.15 (0.10)	0.07 (0.08)	0.11 (0.08)
ATTARI-12		-0.73 (0.72)	-0.69 (0.73)	-0.49 (0.52)	-0.82* (0.37)	-0.73* (0.36)
Trust in AI		0.48 (0.78)	0.42 (0.79)	-0.01 (0.59)	0.51 (0.39)	0.52 (0.36)
Trust in Human		0.42* (0.21)	0.44* (0.22)	0.47*** (0.14)	0.27* (0.11)	0.40*** (0.11)
Positive Reciprocity		0.37 (0.29)	0.44 (0.29)	0.70** (0.24)	0.06 (0.20)	0.13 (0.20)
Age			0.03 (0.04)			0.07*** (0.02)
Male			-0.88 (0.84)			-1.03* (0.44)
NoAI-AI × Others' Contribution				-0.06 (0.07)	-0.15** (0.05)	-0.14** (0.05)
Sample	Full	Full	Full	Full	CC only	CC only
Observations	161	161	158	3,381	2,583	2,520
Adjusted R ²	0.47	0.50	0.50	0.59	0.83	0.83

Notes: The table shows regression estimates on unconditional (Columns 1-3, OLS regressions) and conditional (Columns 4-6, random-effects regressions) contributions in the public goods game. Columns 1-4 report data from the full samples of NoAI-NoAI and NoAI-AI. Columns 5-6 report data from only conditional contribution types in both treatment conditions. Robust standard errors are in parentheses. Levels of significance: *p<0.05, **p<0.01, ***p<0.001.

Table OA-2: Treatment effects in the weakest-link game

	<i>Dependent variable:</i>					
	<i>Unconditional effort provision</i>			<i>Conditional effort provision</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	3.59** (1.13)	-2.18 (3.15)	-7.09* (3.58)	-2.60 (2.26)	-3.20* (1.49)	-4.47** (1.60)
NoAI-AI	-0.40 (0.96)	-0.49 (0.95)	-0.87 (0.95)	-0.22 (0.74)	0.49 (0.33)	0.37 (0.33)
1st-order norm. belief	0.53*** (0.14)	0.47*** (0.13)	0.42*** (0.13)			
2nd-order norm. belief	-0.18 (0.16)	-0.15 (0.15)	-0.05 (0.16)			
Descriptive belief	0.45*** (0.13)	0.44*** (0.12)	0.45*** (0.12)			
Others' Contribution				0.82*** (0.05)	0.97*** (0.02)	0.97*** (0.02)
IOS11		0.07 (0.15)	0.11 (0.14)	0.40*** (0.09)	0.09* (0.04)	0.11* (0.04)
ATTARI-12		-1.47 (0.78)	-1.18 (0.74)	0.12 (0.43)	0.03 (0.25)	0.08 (0.25)
Trust in AI		0.98 (0.77)	1.07 (0.73)	-0.44 (0.48)	0.07 (0.27)	0.10 (0.27)
Trust in Human		0.68*** (0.20)	0.85*** (0.21)	0.14 (0.13)	0.09 (0.06)	0.14* (0.06)
Positive Reciprocity		0.54 (0.33)	0.56 (0.34)	0.19 (0.22)	0.20 (0.14)	0.22 (0.15)
Age			0.10* (0.04)			0.03 (0.02)
Male			-1.64 (0.93)			-0.44 (0.28)
NoAI-AI × Others' Contribution				-0.002 (0.06)	-0.02 (0.03)	-0.02 (0.03)
Sample	Full	Full	Full	Full	CC only	CC only
Observations	161	161	158	3,381	2,898	2,856
Adjusted R ²	0.40	0.46	0.49	0.75	0.93	0.92

Notes: The table shows regression estimates on unconditional (Columns 1-3, OLS regressions) and conditional (Columns 4-6, random-effects regressions) effort provision in the weakest-link game. Columns 1-4 report data from the full samples of NoAI-NoAI and NoAI-AI. Columns 5-6 report data from only conditional coordination types in both treatment conditions. Robust standard errors are in parentheses. Levels of significance: *p<0.05, **p<0.01, ***p<0.001.

OA-3 Experimental Instructions

This section presents an English translation of the main experimental instructions. Screenshots of the original instructions in German are stored at the OSF repository linked to this article. Each line indicates a new screen. Where appropriate, we add comments **** in italic text between asterisks **** above the line. These comments were not shown to participants. Prior to the instructions shown below, participants received information about the study, were asked to provide their informed consent for participation, and confirmed that the participation criteria were met. At the end of the survey, we provided information about the payment.

Part A

In this part of the study, you will make a **cooperation** decision in a group of four people in total: you and three other randomly selected participants.

Each group member receives **20 tokens**. Each member decides for themselves how to allocate their own 20 tokens **between a private account and a group account**. The following applies:

Private account

The amount of the payoff from the private account depends only on **your own** decision. Each token in your own private account increases **your own payoff by 1 taler**.

Group account

The amount of the payoff from the group account depends on the decisions of **all group members**.

Each token in the group account increases the payoff **for each group member by 0.4 talers**.

You are free to decide how many of your 20 tokens you put into your private account and how many you put into the group account. In total, all **20 tokens** must be allocated.

If this part is randomly selected for your group's payment, you will receive a bonus payment corresponding to the number of talers you received from your private account and the group account. Here, **1 taler = €0.10**.

Before you can make your decision for this part, we would like to ask you to answer the following comprehension questions.

How many talers do you receive if you put all 20 tokens into the group account and the 3 other group members also put all 20 tokens into the group account?

- $20 \times 1 = 20$
- $20 \times 0,4 = 8$
- $20 \times 1 + 60 \times 0,4 = 44$
- $80 \times 0,4 = 32$
- $80 \times 1 = 80$

How many talers do you receive if you put all 20 tokens into the private account and the 3 other group members put all 20 tokens into the group account?

- $20 \times 1 = 20$
- $20 \times 0,4 = 8$
- $20 \times 1 + 60 \times 0,4 = 44$
- $80 \times 0,4 = 32$
- $80 \times 1 = 80$

How many talers do you receive if you put all 20 tokens into the private account and the 3 other group members also put all 20 tokens into the private account?

- $20 \times 1 = 20$
- $20 \times 0,4 = 8$
- $20 \times 1 + 60 \times 0,4 = 44$
- $80 \times 0,4 = 32$
- $80 \times 1 = 80$

Before you make your decisions in Part A, we would like to inform you of the following: **You are not allowed to use or consult artificial intelligence to make your decisions in Part A!**

Please note: **The other members of your group have received the same message.**

**** This was shown only in NoAI-NoAI. ****

Before you make your decisions in Part A, we would like to inform you of the following: **You are not allowed to use or consult artificial intelligence to make your decisions in Part A!**

Please note: **Not every member was explicitly prohibited from using artificial intelligence to make decisions in Part A. Therefore, it may be that some of your group members make their decisions with the help of artificial intelligence.**

**** This was shown only in NoAI-AI. ****

Before you make your decisions in Part A, we would like to inform you of the following: **You are allowed to use or consult artificial intelligence to make your decisions in Part A!**

Please note: **Not every member was explicitly allowed to use artificial intelligence to make decisions in Part A. Therefore, it may be that some of your group members make their decisions without the help of artificial intelligence.**

Phase 1

This is **Phase 1** of Part A.

How many of the 20 tokens would you like to put into the group account?

(Note: Tokens that you do not put into the group account will automatically be put into your private account.)

< Slider from 0 to 20. >

Phase 2

This is **Phase 2** of Part A.

You are still in **the same decision situation with the same three group members**. We are now interested in how many **of your 20 tokens you want to put into the group account when you know how the other three group members have decided**.

For this purpose, we present you below with all possible averages, rounded to whole tokens, that the other three group members can put into the group account. If this phase is used for your bonus payment, we will look at how the other three group members decided and match their average to your decision in this phase.

How many tokens would you like to put into the **group account**, given that the other group members put, on average, the number of tokens listed below into the group account?

(Note: Tokens that you do not put into the group account will automatically be put into your private account.)

< 21 sliders varying other's average contribution from 0 to 20. Each slider allows the participant to select between 0 and 20. >

Phase 3

This is **Phase 3** of Part A.

Think briefly about the decision situation in Part A. We will now ask you three questions about

this decision situation. For **Questions 2 and 3, you can receive an additional bonus payment**. To do so, you must identify the average contribution, rounded to whole tokens, that was chosen by most other study participants.

You will receive €0.15 for the correct answer. You will receive €0.14 / €0.12 / €0.08 / €0.03 if your answer differs from the correct answer by 1 / 2 / 3 / 4 tokens, respectively, regardless of the direction. If your answer differs by more than 4 tokens, you will not receive any bonus payment for this question.

Question 1 and Question 2 are about socially appropriate behavior. By this, we mean behavior that you believe most other study participants would consider to be right. In other words: If someone behaves in a socially inappropriate way, other people might be upset about it.

Question 1: What is the minimum number of tokens contributed per group member to the group account that **you** consider socially appropriate?

< *Slider from 0 to 20.* >

Question 2 (Bonus payment): What is the minimum number of tokens contributed per group member to the group account that **most other study participants** consider socially appropriate?

< *Slider from 0 to 20.* >

Question 3 (Bonus payment): What is the number of tokens that most other study participants contributed to the group account?

< *Slider from 0 to 20.* >

Part B

In this part of the study, you will make a **coordination** decision in a group of four people in total: you and three other randomly selected participants.

Each group member has **20 hours of time available**. Each member decides for themselves how to allocate their own 20 hours between a private project and a group project. The following applies:

Private project

The amount of the payoff from the private project depends only on **your own** decision. Each hour in your own private project increases **your own payoff by 1 taler**.

Group project

The amount of the payoff from the group project depends **only on the decision of the group member with the lowest number of hours in the group project.**

Each contributed hour **of this group member (with the lowest number of hours)** in the group project increases the payoff **for each group member by 1.6 talers.**

You are free to decide how many of your 20 hours you contribute to the private project and how many you contribute to the group project. In total, all **20 hours** must be allocated.

If this part is randomly selected for your group's payment, you will receive a bonus payment corresponding to the number of talers you received from your private project and the group project. Here, **1 taler = €0.10.**

Before you can make your decision for this part, we would like to ask you to answer the following comprehension questions.

How many talers do you receive if you contribute all 20 hours to the group project and the 3 other group members also contribute all 20 hours to the group project?

- $20 \times 1 = 20$
- $20 \times 1,6 = 32$
- $20 \times 1 + 60 \times 1,6 = 116$
- $80 \times 1,6 = 128$
- $80 \times 1 = 80$

How many talers do you receive if you contribute 0 hours to the group project (and accordingly contribute 20 hours to your private project) and the 3 other group members contribute all 20 tokens to the group project?

- $20 \times 1 = 20$
- $20 \times 1,6 = 32$
- $20 \times 1 + 60 \times 1,6 = 116$
- $80 \times 1,6 = 128$
- $80 \times 1 = 80$

How many talers do you receive if you contribute 0 hours to the group project (and accordingly contribute 20 hours to your private project) and the 3 other group members also contribute 0 hours to the group project?

- $20 \times 1 = 20$
- $20 \times 1,6 = 32$
- $20 \times 1 + 60 \times 1,6 = 116$
- $80 \times 1,6 = 128$
- $80 \times 1 = 80$

How many talers do you receive if you contribute 4 hours to the group project (and accordingly contribute 16 hours to your private project) and the 3 other group members contribute 2, 5, and 20 hours, respectively, to the group project?

- $4 \times 1,6 + 16 \times 1 = 22,4$
- $20 \times 1,6 + 16 \times 1 = 48$
- $2 \times 1,6 + 16 \times 1 = 19,2$
- $2 \times 1,6 = 3,2$
- $16 \times 1 = 16$

Before you make your decisions in Part B, we would like to inform you of the following: **You are again not allowed to use or consult artificial intelligence to make your decisions in Part B!**

Please note: **The other group members in your group have received the same message.**

**** This was shown only in NoAI-NoAI. ****

Before you make your decisions in Part B, we would like to inform you of the following: **You are again not allowed to use or consult artificial intelligence to make your decisions in Part B!**

Please note: **Not every member was explicitly prohibited from using artificial intelligence to make decisions in Part B. Therefore, it may be that some of your group members make their decisions with the help of artificial intelligence.**

**** This was shown only in NoAI-AI. ****

Before you make your decisions in Part B, we would like to inform you of the following: **You are again allowed to use or consult artificial intelligence to make your decisions in Part B!**

Please note: **Not every member was explicitly allowed to use artificial intelligence to make decisions in Part B. Therefore, it may be that some of your group members make their decisions without the help of artificial intelligence.**

**** This was shown only in AI-AI. ****

Phase 1

This is **Phase 1** of Part B.

How many of the 20 hours would you like to contribute to the group project?

(Note: Tokens that you do not contribute to the group project will automatically be contributed to your private project.)

< Slider from 0 to 20. >

Phase 2

This is **Phase 2** of Part B.

You are still in **the same decision situation with the same three group members**. We are now interested in how many **of your 20 hours you want to contribute to the group project when you know how the other three group members have decided**.

For this purpose, we present you below with all possible values for the lowest number of hours contributed by the other three group members to the group project. If this phase is used for your bonus payment, we will look at how the other three group members decided and match their lowest number of hours to your decision in this phase.

How many hours would you like to contribute to the **group project**, given that the lowest contribution of the other group members to the group project corresponds to the hourly contributions listed below?

(Note: Hours that you do not contribute to the group project will automatically be contributed to your private project.)

< 21 sliders varying the lowest contribution of others' from 0 to 20. Each slider allows the participant to select between 0 and 20. >

Phase 3

This is **Phase 3** of Part B.

Think briefly about the decision situation in Part B. We will now ask you three questions about this decision situation. For **Questions 2 and 3, you can receive an additional bonus payment**. To do so, you must identify the average number of hours, rounded to whole hours, that was chosen by most other study participants.

You will receive €0.15 for the correct answer. You will receive €0.14 / €0.12 / €0.08 / €0.03 if your answer differs from the correct answer by 1 / 2 / 3 / 4 hours, respectively, regardless of the direction. If your answer differs by more than 4 hours, you will not receive any bonus payment for this question.

Question 1 and Question 2 are about socially appropriate behavior. By this, we mean behavior that you believe most other study participants would consider to be right. In other words: If someone behaves in a socially inappropriate way, other people might be upset about it.

Question 1: What is the minimum number of hours contributed per group member to the group project that **you** consider socially appropriate?

< Slider from 0 to 20. >

Question 2 (Bonus payment): What is the minimum number of hours contributed per group member to the group project that **most other study participants** consider socially appropriate?

< Slider from 0 to 20. >

Question 3 (Bonus payment): What is the number of hours that most other study participants contributed to the group project?

< Slider from 0 to 20. >

Questionnaire

As soon as you move the slider below, two circles will appear in the field. The position of the slider determines how strongly the circles overlap. When the slider is all the way to the left, the circles look like this < *image where sliders are not overlapping* >. When the slider is approximately in the middle, the circles look like this < *image where sliders are partly overlapping* >. When it is all the way to the right, the circles look like this < *image where sliders are almost completely overlapping* >.

You should interpret the degree of overlap as a representation of the relationship between you, represented by “You”, and “X”. **Here, “X” refers to the other study participants in your group who were also not allowed to use artificial intelligence to make the decisions.**

Please position the slider so that the circles indicate the extent to which you and “X” are connected.

*** *This was shown only in NoAI-NoAI.* ***

Questionnaire

As soon as you move the slider below, two circles will appear in the field. The position of the slider determines how strongly the circles overlap. When the slider is all the way to the left, the circles look like this < *image where sliders are not overlapping* >. When the slider is approximately in the middle, the circles look like this < *image where sliders are partly overlapping* >. When it is all the way to the right, the circles look like this < *image where sliders are almost completely overlapping* >.

You should interpret the degree of overlap as a representation of the relationship between you, represented by “You”, and “X”. **Here, “X” refers to the other study participants in your group who were allowed to use artificial intelligence to make the decisions.**

Please position the slider so that the circles indicate the extent to which you and “X” are connected.

**** This was shown only in NoAI-AI. ****

Questionnaire

As soon as you move the slider below, two circles will appear in the field. The position of the slider determines how strongly the circles overlap. When the slider is all the way to the left, the circles look like this < *image where sliders are not overlapping* >. When the slider is approximately in the middle, the circles look like this < *image where sliders are partly overlapping* >. When it is all the way to the right, the circles look like this < *image where sliders are almost completely overlapping* >.

You should interpret the degree of overlap as a representation of the relationship between you, represented by “You”, and “X”. **Here, “X” refers to the other study participants in your group who were also allowed to use artificial intelligence to make the decisions.**

Please position the slider so that the circles indicate the extent to which you and “X” are connected.

**** This was shown only in AI-AI. ****

In the following, we are interested in your attitudes towards artificial intelligence (AI). AI can execute tasks that typically require human intelligence. It enables machines to sense, act, learn, and adapt in an autonomous, human-like way. AI may be part of a computer or online platform - but it can also be encountered in various other hardware devices such as robots.

Please indicate how strongly you agree or disagree with the following statements:

AI will make this world a better place.

- Strongly disagree
- Somewhat disagree
- Neither agree nor disagree
- Somewhat agree
- Strongly agree

I have strong negative emotions about AI.

- Strongly disagree
- Somewhat disagree
- Neither agree nor disagree
- Somewhat agree
- Strongly agree

I want to use technologies that rely on AI.

- Strongly disagree
- Somewhat disagree
- Neither agree nor disagree
- Somewhat agree
- Strongly agree

AI has more disadvantages than advantages.

- Strongly disagree
- Somewhat disagree
- Neither agree nor disagree
- Somewhat agree
- Strongly agree

I look forward to future AI developments.

- Strongly disagree
- Somewhat disagree
- Neither agree nor disagree
- Somewhat agree
- Strongly agree

AI offers solutions to many world problems.

- Strongly disagree
- Somewhat disagree
- Neither agree nor disagree
- Somewhat agree
- Strongly agree

I prefer technologies that do not feature AI.

- Strongly disagree
- Somewhat disagree
- Neither agree nor disagree
- Somewhat agree
- Strongly agree

I am afraid of AI.

- Strongly disagree
- Somewhat disagree
- Neither agree nor disagree
- Somewhat agree
- Strongly agree

I would rather choose a technology with AI than one without it.

- Strongly disagree
- Somewhat disagree
- Neither agree nor disagree
- Somewhat agree
- Strongly agree

AI creates problems rather than solving them.

- Strongly disagree
- Somewhat disagree
- Neither agree nor disagree
- Somewhat agree
- Strongly agree

When I think about AI, I have mostly positive feelings.

- Strongly disagree
- Somewhat disagree
- Neither agree nor disagree
- Somewhat agree
- Strongly agree

I would rather avoid technologies that are based on AI.

- Strongly disagree
- Somewhat disagree
- Neither agree nor disagree
- Somewhat agree
- Strongly agree

I generally trust decisions made by AI systems.

- Strongly disagree
- Somewhat disagree
- Neither agree nor disagree
- Somewhat agree
- Strongly agree

I usually trust a human decision more than an AI decision, even when the AI is said to be very accurate.

- Strongly disagree
- Somewhat disagree
- Neither agree nor disagree
- Somewhat agree
- Strongly agree

I am willing to rely on AI systems when I have to make difficult decisions.

- Strongly disagree
- Somewhat disagree
- Neither agree nor disagree
- Somewhat agree
- Strongly agree

I am skeptical of decisions made by AI systems.

- Strongly disagree
- Somewhat disagree
- Neither agree nor disagree
- Somewhat agree
- Strongly agree

I feel at ease when AI systems are used to support decisions in everyday life.

- Strongly disagree
 - Somewhat disagree
 - Neither agree nor disagree
 - Somewhat agree
 - Strongly agree
-

How well does each of the following statements describe you as a person? 0 means “does not describe me at all” and 10 means “describes me perfectly”.

If someone does me a favor, I am willing to return it.

< *Slider from 0 to 10.* >

I assume that people only have the best intentions.

< *Slider from 0 to 10.* >

Which gender do you identify with?

- Female
- Male
- Other
- No response

How old are you? < *Integer field.* >
