Historical Social Research

Maximilian Noichl & Johannes Marx

Simulation of Group Agency – From Collective Intentions to Proto-Collective Actors

doi: 10.12759/hsr.48.2023.25

Published in:
Historical Social Research 48 (2023) 3

Cite as:

Noichl, Maximilian & Johannes Marx. 2023.
Simulation of Group Agency –
From Collective Intentions to Proto-Collective Actors.
Historical Social Research 48 (3): 56-74. doi: 10.12759/hsr.48.2023.25

Simulation of Group Agency – From Collective Intentions to Proto-Collective Actors

Maximilian Noichl & Johannes Marx *

Abstract: *»Kollektive Intentionen und proto-kollektive Akteure«.* This paper investigates the conditions under which cooperative team reasoning arises and stabilizes in complex social structures. Team reasoning is a theory that explains cooperative behavior in social settings of strategic choice, even in situations where classical game theory fails. By simulating the emergence of cooperation via team reasoning, this paper analyzes the performance of team reasoners compared to classically rational agents and individual reasoners. Simulation results are provided regarding the efficacy of team reasoning in mixed-game settings. It is shown that cooperative team reasoning is viable and stabilizing under favorable conditions such as the share of coordination games played, but sensitive to the amount and abilities of their interacting counterparts. Finally, the paper provides first ideas on how the current framework might be extended toward collective actors that gain further stability through processes of self-formalization and inner-organizational redistribution.

Keywords: Team reasoning, collective agency, agent-based modelling, computer simulation, game theory.

1. Introduction

Team reasoning is a theory to explain cooperative behavior in social settings of strategic choice, even in situations where classical game theory fails. It takes its theoretical starting point in the well-documented discrepancy between, on the one hand, the empirically observed behaviors and judgments

* Maximilian Noichl, Political Science Department, University of Bamberg, Feldkirchenstraße 21, 96052 Bamberg, Germany; noichlmax@hotmail.co.uk. ORCID iD: 0000-0003-4518-0837.
Johannes Marx, Political Science Department, University of Bamberg, Feldkirchenstraße 21, 96052 Bamberg, Germany; johannes.marx@uni-bamberg.de. ORCID iD: 0000-0003-3090-4097.
Acknowledgements: We would like to thank Leyla Ade and Olivier Roy for their extensive comments and discussions, which helped to develop the ideas presented in this paper. Additionally, we are grateful to the participants of the 2022 workshop on Team Reasoning and Collective Agency, as well as the audiences in Helsinki, Bamberg, and Bochum for their insightful feedback. This project was funded by the German Research Foundation (DFG, project number 443226175).

of experimental subjects exposed to simple one-shot games and, on the other, the results that classical game theory would suggest for them (Bacharach 1999, 2006; Colman, Pulford, and Lawrence 2014). Generally, the ability of real actors to coordinate themselves seems to exceed the coordination predicted in classical game theory like Stag Hunt or HiLo games by far. The theory of *circumspect team reasoning* (Bacharach 2006) extends classical game theory through a novel system for the determination of equilibria while making no drastic changes to its classical game theoretic foundations and without introducing additional axioms. To model team reasoning, Bacharach introduces the idea of a fictional coordinator (the "team planner") that lets rational agents coordinate their actions to strategy profiles. In combination with an individual estimation of team reasoning probabilities by the actors, the theory of Bacharach promises to provide solutions to games that both match our empirical observations and our theoretical demands for a solution in line with the standard rationality assumptions of game theory. Simply put, an actor engaged in team reasoning is not asking, *what should I do* in a particular situation, but *what should we do in that situation and what is my part in making that happen*. Addressing this question enables rational actors to solve coordination games and even social action dilemmas that are characterized by a tension between self-interests and a Pareto-efficient collective output.

While team reasoning is well understood for classic games in simple settings, the conditions under which it arises and stabilizes in more complex social structures remain largely unexplored. Analytical solutions to these questions are mainly focusing on simple interaction structures and concentrate primarily on single game-theoretic settings in the body of literature. There is some literature discussing the effects of the ludic ecology on the stability of team reasoning in evolutionary settings (for an overview, see Ade and Roy 2023, in this special issue). It is at least questionable to what extent analytic solutions can be found at all due to the complexity of the interaction situation in large, connected groups playing a mix of games. At the same time, simulational results on the emergence of stable team reasoning are rare and show no clear picture. If team reasoning is able to explain cooperation with fewer assumptions and in harder cases than alternative theories, it should be considered as a main driver for collective agency. Thus, the explanation of such complex collaborative situations presents a clear desideratum for political and social theory (Elsenbroich and Payette 2020; Newton 2017). To address this lacuna, we develop a model for the simulation of the emergence of team reasoning to analyze conditions for its stability, which allows us to compare the performance of team reasoners with that of classically rational agents and individual reasoners (see section 3 for a description of the characteristics of the agents).

In the present contribution, we will accordingly give a sketch of the theoretical research on team reasoning, introduce the simulational framework, argue for a series of theoretical choices that we have made in its set-up, and apply it to several game-theoretical standard situations. We find team reasoning to be somewhat viable and stabilizing under slightly favorable conditions, in line with some earlier literature (Angus and Newton 2015; Elsenbroich and Payette 2020), but sensitive to the amount and abilities of classical rational agents, as suggested by Paternotte (2018). We further provide simulation results regarding the efficacy of team reasoning in several mixed-game settings. Finally, we give some indications on how the current framework might be extended toward collective actors that gain further stability through self-formalization and redistribution. In doing so, we aim to contribute to shedding light on the role of team reasoning to the emergence of collective actors.

2. Team Reasoning and Open Questions

One of the basic methodological assumptions in most parts of social science is that agency is located in individual actors (methodological individualism). This assumption includes that (i) institutionalized social groups are not treated as actors of their own right and that (ii) actors perform actions in line with their own intentions. Condition (i) demands that speaking of the behavior of states, firms, organizations, or groups of actors can only be understood as a shortcut for a fully fleshed-out explanation. Such an explanation should explain the "behavior" of such collective bodies referring to the preferences and beliefs of the individual agents who are members of that social entity. Condition (ii) requires that actors behave on their own preferences and beliefs, i.e., it should be the preferences and beliefs of the actor herself who govern her actions. This condition does not exclude that the actors' preferences are social or altruistic and refer to others.

Team reasoning challenges both conditions.[1] It allows for teams of agents to be understood as agents in their own right, and it relaxes the conditions that the (originally given) actor's own preferences and beliefs should drive her actions.[2] Team reasoning takes its origins from the observation that there are a number of game-theoretic puzzles in which conventional game theory both fails to predict the behavior of human players and even might predict things most humans find highly suspect. In the HiLo game (Fig. 4b), for example, classical game theory finds two Nash equilibria – the reception of a payoff is after all conditional on the other player choosing the same option – but it seems rather intuitive that the option that promises the higher utility is

[1] Compare Radzvilas and Korpus (2021), who argue that team reasoning is compatible with ontological individualism.

[2] However, the actor will still act on her preferences, but the team preference will become her preference and the utility she is striving for is the team utility. Only in this weak sense, the actor is not acting on behalf of her "own" preferences.

the superior one. And unsurprisingly, experimental data shows that players can successfully coordinate their actions to Hi in the HiLo game. Furthermore, it is a well-documented observation that the proportion of players choosing cooperation is above 40 percent in a one-shot Prisoner's Dilemma (Andreoni and Miller 1993; Camerer 2003; Heuer and Orland 2019).

These empirical results would not be a problem for a theory of rational action as long as the irrationality of the actors could explain that discrepancy. However, following Gold, it is "at least arguably rational" (Gold and Sugden 2007, 117) to follow strategies that can lead to Pareto superior outcomes for the group of players. Of course, one can argue that every player performs individually better in a Prisoner's Dilemma game if she defects. That is the reason for characterizing the strategy as dominant. In that sense, "defection" would answer the reasoning process of a rational actor asking what she should do to perform best. However, with the same line of reasoning, one could argue that still, both players would perform better if they both cooperate. If a player realizes this and stops asking what she should do to succeed and starts asking what they as a team should do to perform best, then mutual cooperation becomes admissible. Following this line of argumentation, it would be rational for a person to reason that cooperation would be the best thing to do. This change in the reference point for the actors' reasoning process is the core of the theory of team reasoning.

The basic idea behind team reasoning is that a lot of simple coordination and strategic interaction problems can be solved easily enough by introducing a central coordinator who provides a team preference by ordering all outcomes. This allows the team reasoners to determine the best outcome for the group (first rank of the team preference). This fictitious coordinator then can assign every agent the task they have to accomplish in order to obtain this result and the team members adopt the team preference as their own. This fictitious coordinator represents in the game theoretic model the perspective of team reasoners on the game.

One technical way to implement this line of reasoning in game theory is to model the team as a player of its own right in the game. Bacharach's version of team reasoning, which we follow to a large extent in this contribution (Bacharach 1999), formally implements this idea by explicitly adding an additional player (the "Team") to each game. The payoff of the team as a player is determined by the average of the individual players' payoffs. In its role as central coordinator, the "team" does not select individual actions but action profiles for all players – in the case of the Prisoner's Dilemma for example, the combinations $[CC, CD, DC, DD]$. A 2×2 game thus gets transformed into a $2 \times 2 \times 4$ three-player game.

Figure 1 Illustration of the Team Equilibrium

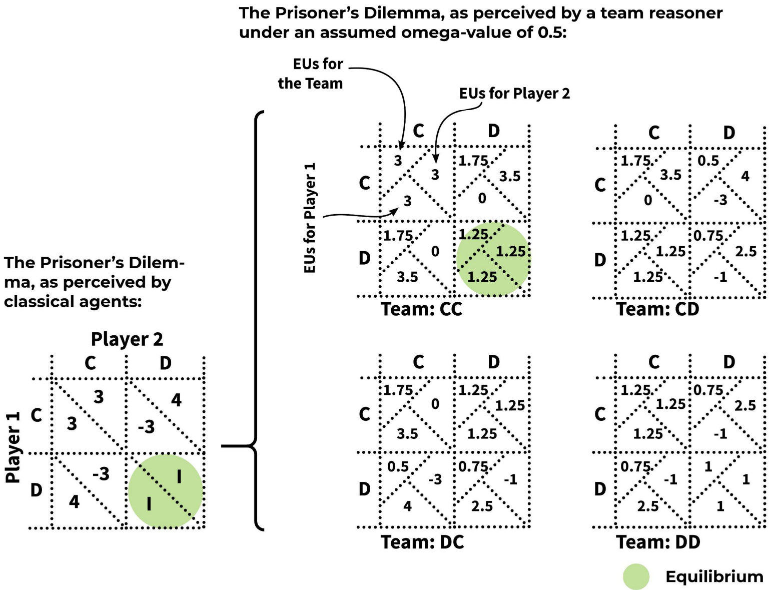

The two-player game gets transformed into a three-player game which includes the team. EU stands for expected utilities. In our case, the Nash-equilibrium in the large game lies in the CC-assignment for the team, which therefore prescribes the strategy C to the team reasoner. This of course changes when omega falls below 0.333.

Team reasoning now consists in finding the Nash equilibria in this novel game and in playing the team profile selected by them. One should note that the choice of the average of what the team members consider as team-payoff for the function generating the team utilities is not equivocal, and many other schemes could be considered (see Karpus and Gold 2016; Sugden 2000, 2015). However, this has no bearing on the actual payoffs they receive from the games they play. The actual payoffs which they receive in the simulation are drawn from the same games for individual and team reasoners.

Of course, there are no such teams as real players alongside the individual players in situations of strategic choice. Rather the team player represents a particular perspective that individual players can follow. Bacharach and Stahl (2000) refer to the concept of frame theory to describe the transformational processes individual actors experience when they shift to this perspective. Actors who behave in line with the we-frame no longer ask what is best for them individually but instead what is the best result we as a team can receive. This shift includes an agency transformation – i.e., the agents identify with the team and there is common knowledge to what degree all actors identify with the team – and a payoff transformation – i.e., the payoff does not represent the individual utility classical agents receive but the payoffs for the team. In this perspective, both of the basic assumptions of methodological individualism are challenged. There are various reasons why actors can switch from I-mode frames to we-mode frames, ranging from psychological processes to game-induced effects (in pure coordination games with one Pareto efficient equilibrium).[3]

Team reasoning seems to be useful to enable actors to cooperate in settings of strategic choice. However, team reasoners are rational, e.g., their willingness to cooperate is not unconditional. Cooperation is risky and unlikely to occur as long as the individual agent perceives it as unlikely that their co-players will perform the same transformation to arrive at the best team outcome. To reflect this uncertainty, a parameter ω, is introduced in the theory of team reasoning. In our simulation ω reflects for each agent their estimate of the probability that they are playing with another team reasoner. In the $2 \times 2 \times 4$-matrix, all utilities are then replaced by expected utilities, given a specific ω. This theoretical move implies that in many games, for low values of ω, the team equilibrium can remain in a non-cooperative state – team reasoning does not exhaust itself in mere altruistic cooperation. In Fig. 4 we show the relationship between different values for ω, and the achievement of team-play in more detail.

We thus suggest that team reasoning might be a suitable theory of cooperation in a number of domains, as it satisfies a demand for rational cooperation in one- or few-shot interactions and does so in an empirically plausible fashion. But it is unclear whether this theoretical expectation really holds in simulations in a changing ludic environment with different shares of team reasoners and classical agents. The results in the research literature have so far been mixed in this regard (see Ade and Roy 2023, in this special issue for an overview). Amadae and Lempert (2015), for example, show that in an evolutionary setting with agents playing a Prisoner's Dilemma, team reasoning is inferior to other reasoning modes and becomes extinct. While on the other side in a slightly different setting, Newton (2017) finds evidence that team reasoning seems to be stable even in a Prisoner's Dilemma setup. Furthermore, there is some evidence that team reasoning is evolutionarily stable in HiLo games. Nevertheless, it remains unclear to what extent the results are robust if the agents have to choose their interacting partners in different neighborhoods. Finally, the extent to which team reasoning leads to stable cooperation in situations characterized by a mix of game types has mainly been investigated by means of evolutionary game theory. We are particularly interested in whether these results can be replicated in computer simulations and to what extent the results are stable if we make classical rational agents slightly

3 For example, see Gold and Sugden (2007) and Gold (2017) for different accounts of team reasoning.

more realistic and go beyond the fiction of the conventional homo economicus.

3. A Computational Framework to Model Team Reasoning

To explore the theory of team reasoning, we develop a computational framework that implements Bacharach's basic model of unreliable team interactions and integrates it into a full simulation environment consisting of multiple types of actors, learning mechanisms, games, and social structures. The whole framework is implemented in Python and is designed to be easily adapted to different environments.[4] This allows us to apply it to a variety of theoretical situations, as well as to fit it to empirically gathered parameters. Below we document the individual components of each simulation.

3.1 Actors

A central feature of Bacharach's theory is to explain the situation in which agents need to coordinate with others who might or might not team reason. Therefore, in our simulations we distinguish between team reasoning agents (team reasoners), and agents that play strategies that would be considered optimal in conventional game theory (classical agents), i.e., if available they play dominant strategies or mixed strategies in games as HiLo. Team reasoners are modeled as actors who maximize expected utilities. Later we will introduce an additional mixed-type, learning classical agents (individual reasoners). In the presented work, we run simulations consisting of 100 agents, which are placed on a von Neumann grid, as illustrated in Fig. 2. This scale allows for the emergence of small to medium-sized proto-collective actors, which could correspond to working groups, bands of hunter-gatherers, parliaments, international organizations, etc.

We systematically vary the number of team reasoners to investigate at what point cooperative team reasoning becomes stable and more worthwhile than the standard game theory approach of classical agents.

[4] This model is integrated in one software package. It can be used for the simulation of team reasoning in complex social settings while allowing for variation, for example, in the games played by the agents, the number of agents, and the proportion of team reasoners and classically rational agents.

Figure 2 A Simulation Running on a Grid-Model (implemented in NetworkX, Hagberg et al. 2008)

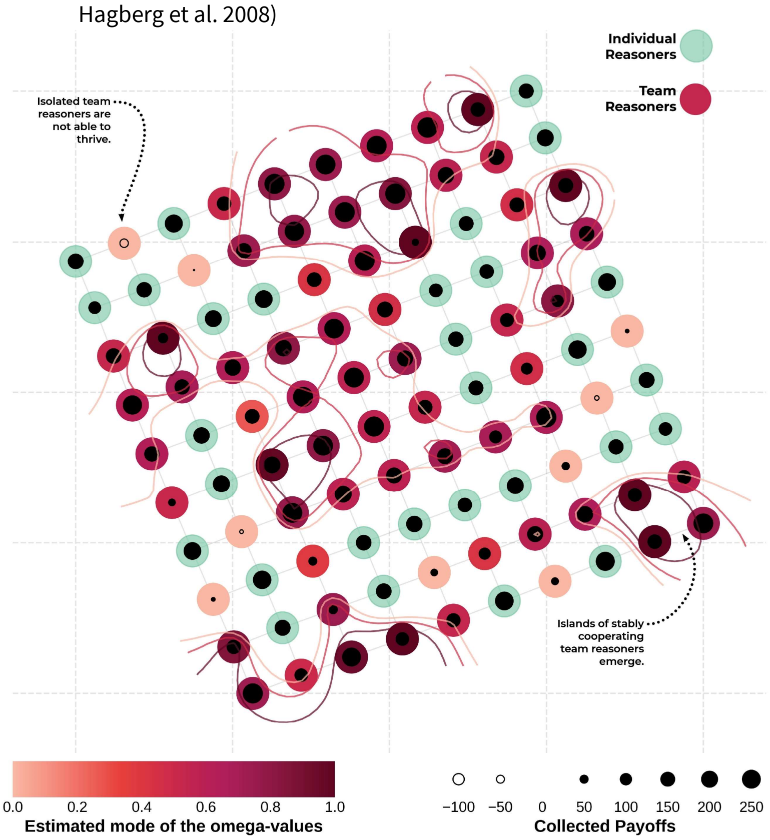

Each dot represents an agent, and their links represent potential interaction partners. We show several emergent clusters of team reasoners forming proto-collective actors visually tied together by contour lines, which are derived by interpolation of the omega values along the grid. In this particular simulation run, 70% of the agents were team reasoners, 30% were classically rational agents. The game-mixture consisted of 50 % HiLo, 50 % Prisoner's Dilemma.

3.2 Learning

The two types of actors differ in the baseline simulation with respect to their ability to learn. Classical reasoners play the games as one-shot games. They have no information about the actors around them and the strategies played so far. Team reasoners, on the other hand, try to use this information. As we have discussed earlier, team reasoning agents have to learn an estimate of the team reasoning probability, ω, of their co-players. Determining ω is quite a tricky problem. After all, ω is not simply a cooperation probability, because team reasoning co-players might arrive at a non-cooperative outcome while team reasoning. Also, as Fig. 4 makes quite transparent, team reasoning agents ought to learn different things from different games, as the threshold for cooperative or coordinated behavior both differ and signify different things: From the payoff in a single HiLo-game, an agent can, for example, learn nothing about the identity of their co-players, as every payoff could also be achieved through the actions of a classical agent playing a mixed strategy. In the Prisoner's Dilemma, on the other hand, cooperation is an unambiguous sign of the presence of a team reasoner. Importantly, although agents are placed at fixed points of their grids, and only play with their direct neighbors for the whole run of each simulation, they are not informed about which of their neighbors they played which round. This limits their ability to engage in tit-for-tat, or similar strategies, which would undercut the situations that necessitate team reasoning in the first place.

To solve this problem, in our proposed framework we let agents directly optimize their payoffs. To do so, we model them as a function of the ω-values through a Gaussian Process that has both a time component, which detects shifting trends, and synchronic functions that relate current omega-values to their expected payoffs. The agents build, maintain, and exploit this model through a process of adaptive Bayesian optimization, which we have adapted from Nyikosa, Osborne, and Roberts (2018).[5] Adaptive Bayesian optimization has been proposed as a method for the estimation of the minima of temporally shifting objective functions with costly evaluation.[6]

At each step, team reasoners can decide to either explore a new ω-value, or to exploit what they have so far learned as the current optimum value for a maximal payoff. To ensure that agents are able to dynamically react to changes in their environment, we also introduce a window parameter (usually set to 12 steps). Only results that lie within the window are considered for the model-fit and in the optimization process. If, for example, an agent has made the experience that a high ω-value is not worth the risk in the current environment, they might at a later point re-examine that decision. A demonstration of how this can look like is given in Fig. 3.

It is important to note that agents here are neither trying to approximate the true relationship between estimates of the amount of team reasoners and payoffs, nor are they necessarily trying to get at the true value of ω. As learning about the parts of the distribution which are less than optimal is costly (as the agents might get dealt the sucker payoff in a Prisoner's Dilemma), they avoid doing so as much as possible and are in some situations quite content with a picture that is incomplete but good enough for their decision-making. Further, this implies that agents can in principle give up their best estimates in favor of inexact but productive ones.

---

5 See also Snoek, Larochelle, and Adams (2012).

6 It has been suggested that Bayesian optimization also yields an empirically adequate account of human exploration dynamics (Borji and Itti 2013).

Figure 3 Demonstration of How Actors May Evaluate Risk

a) Example of the learning process of an individual agent

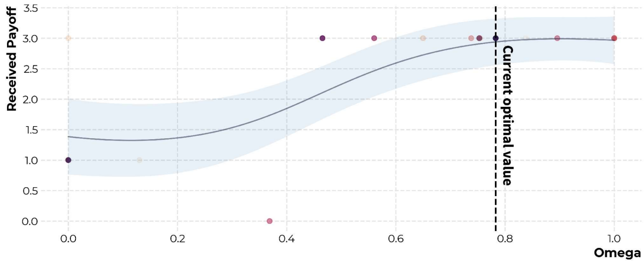

b) Aggregated learning functions of all team reasoners in one simulation

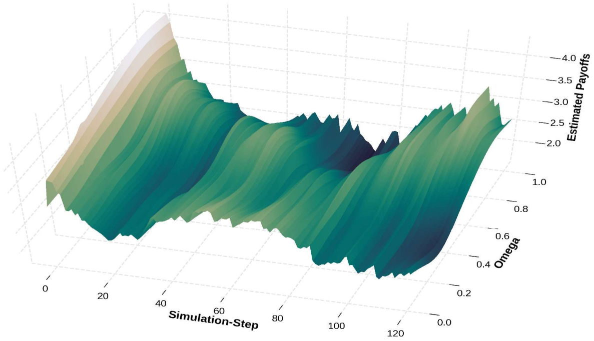

c) Best predictions of all team reasoners in one simulation

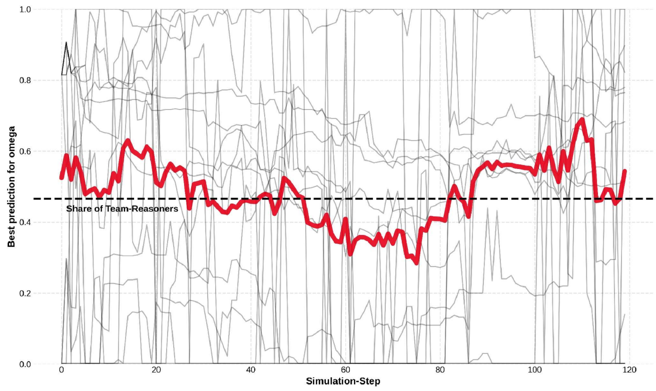

(a) A single agent's determination after 30 iterations of the simulation. After exploring several low values, in this simulation, the agent learned to expect high levels of cooperation, and now mainly exploits high ω-values. (b) The aggregated models behind the learning mechanisms of all team reasoning agents in one example run of the simulation. In the presented case, the agents quickly settle on a medium value for ω, around which they oscillate while considering both higher and lower alternatives collectively at different stages. This pattern becomes more visible in the aggregated predictions for the optimal ω-values the agents hold over time (c).

3.3 Games

Our current framework allows for the in-cooperation of arbitrary 2×2 2-player-games. The behavior of classical agents is modeled in line with orthodox game theory, i.e., classical agents will choose the dominant strategy as available or play mixed strategies in HiLo and Stag Hunt games. The case is different for team reasoners who maximize expected utilities. Because the computation of Nash equilibria in the altered games at different ω-values is computationally expensive, we precompute a number of solutions (represented by the dots in Fig. 4), making use of NashPy (Knight and Campbell 2018) and Gambit (McKelvey, McLennan, and Turocy 2006). As team reasoners do not need to know the exactly expected utilities associated with their current ω-value, but only which profile a certain value selects, it suffices to check at each step whether the closest precomputed ω-values agree on the course of action. Only in the edge cases in which they do not, like, e.g., around ⅓ in the Prisoner's Dilemma, we have to compute an additional Nash equilibrium at runtime.

In the reported simulations, agents encounter the three seminal games HiLo, Stag Hunt, and the Prisoner's Dilemma in a variety of ratios. The switching points, at which team reasoners begin to cooperate in the Prisoner's Dilemma or eliminate an equilibrium in HiLo or Stag Hunt are reported in Fig. 4.

Figure 4 The Expected Utilities or a Team Reasoner at the Equilibria in the Prisoner's Dilemma (a), HiLo (b), and Stag Hunt (c)

a) Prisoner's Dilemma

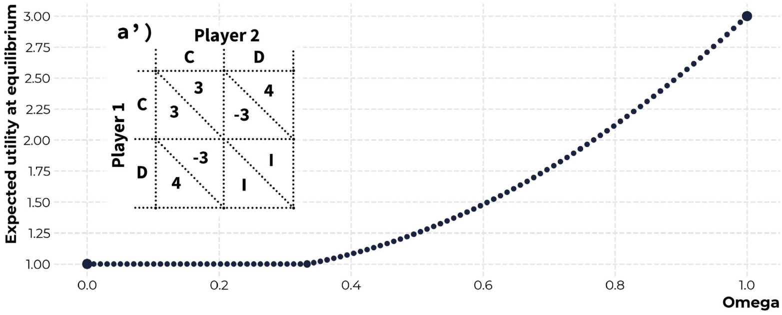

b) HiLo

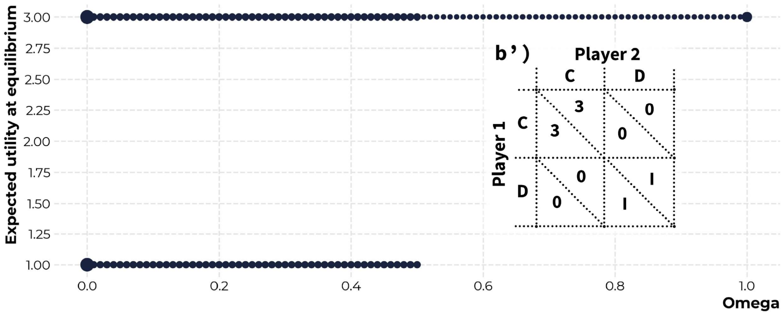

c) Stag Hunt

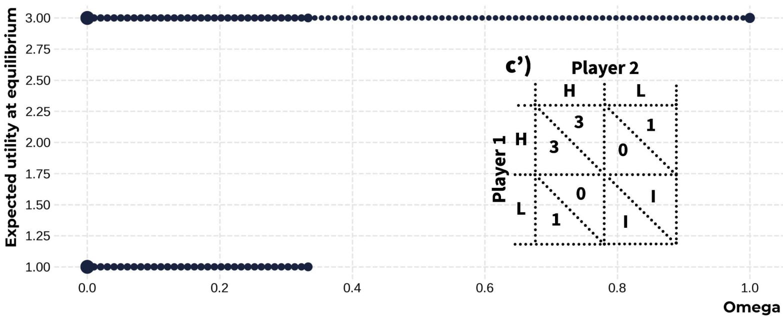

In the exemplar Prisoner's Dilemma, the team reasoning agent switches to cooperation at $\omega = 0.3$. In HiLo they play a mixed strategy until an ω of 0.5, when settling on "high" becomes more effective. Below that value, the chance of being paired with an agent that does not team reason, and thus plays randomly, is too high. The same is true for Stag Hunt, but here the agents coordinate already at a lower threshold.

4. The Conditions for Stable Teams

The results of a range of simulations run under these settings are illustrated in Fig. 5.

Figure 5 Results of a Range of Simulations in which Team Reasoners Play on a Grid with Classical Agents

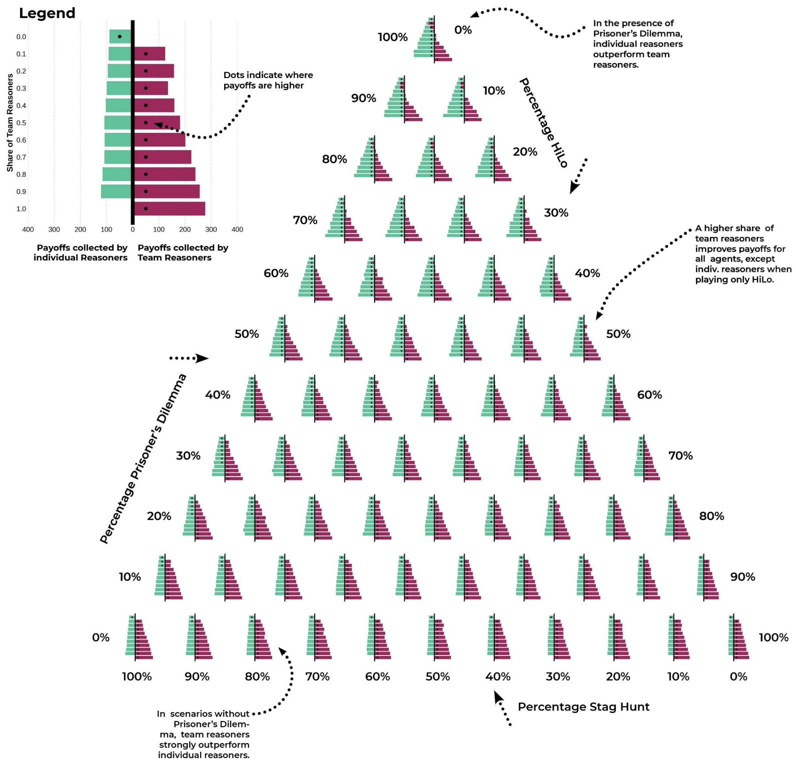

The rows and diagonals of the large triangle represent the mixture of games that are played in the simulations, with the corners identifying simulations in which only one type of game was played. Each small bar chart reports the mean payoffs of team reasoners (violet, right) and classical agents (turquoise, left), under various shares of team reasoners, in each small graphic from top to bottom. We note that in all simulations (except the one in which only HiLo is played) a higher share of team reasoners is beneficial to all players. Further, a higher share of Prisoner's Dilemma among the played games reduces team reasoners' payoffs in comparison to that of classical agents, while a higher share of Stag Hunt and HiLo increases it.

We note that a higher probability to play Prisoner's Dilemma, as opposed to HiLo or Stag Hunt games, favors the classical agents in terms of received payoff, as they, in contrast to the team reasoners, never expose themselves to the possibility of being defected on. The team reasoners' advantage, on the other hand, mainly seems to lie in the exploitation of conditions in which HiLo makes coordination of efforts necessary. This configuration is not surprising: Even when team reasoners find themselves in a situation in which stable cooperation in the Prisoner's Dilemma becomes possible, non-team reasoners will reliably gain by defecting on them. The exact points at which the payoffs of both groups reach parity, of course, hinges on the payoff structures of the games. A Prisoner's Dilemma with a more drastic "sucker-payoff" will increase the burden on team reasoners, while HiLo and Stag Hunt games with a higher reward for coordination will serve to jump-start more cooperation. Overall, the total sum of all payoffs rises the more team reasoners are present in the simulation. Also, team reasoners always do better the more of them are present in the simulation. Non-team reasoners, correspondingly, also tend to do better in scenarios in which more team reasoners are present, as long as the agents do not purely play HiLo, profiting from the cooperation benefits in Stag Hunt and from defecting on team reasoners in the Prisoner's Dilemma. These results are largely in line with those of Amadae and Lempert (2015).

5. A Complication: Learning Individual Reasoners

However, the way we have set up the previous simulation does seem to slightly favor the team reasoners: While they are not able to punish individual defectors or engage in complex communication through games, they are able to trace the general conditions and adapt their strategies to them – something that individual reasoners arguably also ought to be able to do. In the baseline simulation, the classical agents did not use that kind of information. Conveniently, team reasoning also gives us an apparatus to calculate the actions of individual reasoners who are aware that team reasoners are present in their population, namely by simply constructing the three-player game of the team and the two individuals and then not choosing the team's strategy, but their own. For the Prisoner's Dilemma, this leaves things unchanged, as it remains rational for individual reasoners to defect even, especially, in the presence of team reasoners. In HiLo and Stag Hunt games, on the other hand, the play of the individual reasoners becomes identical to team reasoners. We have produced the results of a range of simulations under these assumptions in Fig. 6. The learning individual reasoners are now imbued with an ω-value in the same way the team reasoners are; they just differ in the conclusions they draw from it. We note that the advantage of team reasoners has disappeared under these conditions. The effect, that while at least some Prisoner's Dilemmas are present, a higher share of team reasoners leads to higher average payoffs for all agents does remain though. This suggests that in an evolutionary investigation of team reasoning close attention should be paid to the mechanism of selection. Those selection mechanisms that reward the payoff achieved relative to co-players will likely favor individual reasoners, while threshold-based selection will, under some configurations, favor team reasoners as the global payoff that is achieved commonly rises with the share of team reasoners, which makes it more likely that selection thresholds are met.

Figure 6 Results of a Range of Simulations in which Team Reasoners Play with More Sophisticated Individual Reasoners

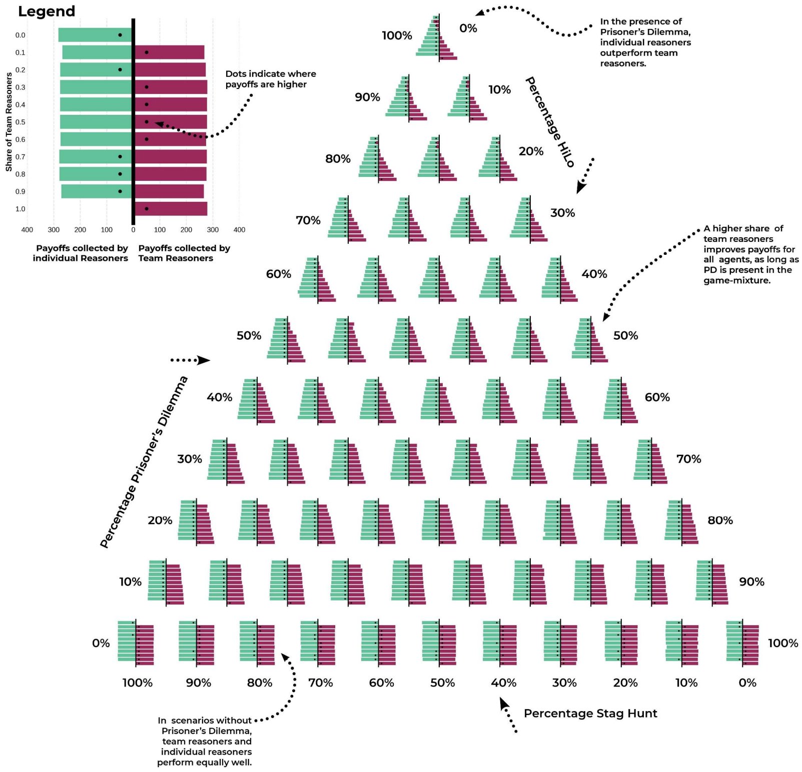

These individual reasoners are able to solve coordination problems and learn about the presence of team reasoners. Here, the advantage of team reasoners vanishes.

6. Prospects: From Stable Teams to Proto-Collective Actors

With the previous simulations, we can show under which conditions teams of actors that are willing to cooperate emerge and develop stable we-intentions. Now, we go one step further and discuss under which conditions groups of actors who have stable we-intentions may be understood as collective actors in a more "formal" sense. We thus link to the sociological literature in the tradition of Coleman (1974; see also Vanberg 1982), which understands collective actors as organizations (see Gehring and Marx 2023, in this special issue).

There are several lines of demarcation proposed between collective actors and mere collections. One of the most general approaches might be the one that treats collective agents as analogous – and therefore committed to the same definitorial demands – to individuals in the biological realm, which have to be distinguished from mere environmental occurrences (Lewis-Martin 2022). This view suggests that actors ought to exist over extended periods of time and are to be understood as self-stabilizing systems (see Meincke 2019). These authors stress that biological individuals, and by extension collective actors, can be understood as operationally closed systems: "Thus, if process A sustains process B, which sustains process C, which sustains process A, then the system ABC is operationally closed" (Lewis-Martin 2022, 283). Importantly, these processes are precarious, which here means that these processes are essential for the survival of the system: "in the absence of the enabling relations established by the operationally closed network, a process belonging to the network will stop or run down" (Paolo and Thompson 2014, 72). At the same time, collective actors are characterized by the fact that these mechanisms contribute to the stabilization of the system, even when individual parts of the system are replaced, cease to exist, or become dysfunctional. Even different configurations at the micro level with different individual agents as members are possible while maintaining the collective actor. In this sense, the collective actors supervene over groups of stably cooperating team reasoners (see Gehring 2023, in this special issue).

Our simulations currently seem to exhibit exactly this kind of self-stabilizing behavior of team reasoners. When a set of agents interacts persistently, they mutually stabilize their expectations that cooperation is worthwhile (we illustrate this in Fig. 2). Even single agents of that group can drop out while the system as a whole will be maintained. Only if these confirmations start to systematically fall away will the system become endangered.

The patterns in our simulations therefore can be sensibly considered to model the minimal conditions for collective actors. However, collective actors in the social world typically exhibit further characteristics than the emergence of collective intentions (see Gehring and Marx 2023; Gilbert 2023, both in this special issue). The collective actors considered by political science and economics, for example, are typically formally constituted. They often come into being through a deliberate, rational decision. One way this can happen is through the transfer of resources. This can include not only material resources but also decision-making powers that are transferred to the collective actor to be created. In social science terms, such actors can be understood as organizations. In our present simulations, we do not model such processes of the emergence of formally organized collective agents. Implementing transfer mechanisms within self-determined groupings between agents, and observing the interaction of these mechanisms with team reasoning, will be one of the central forthcoming challenges.

7. Conclusion

To summarize, we have argued that the theory of team reasoning merits more extensive consideration, as it (1) solves several important problems in conventional game theory in an elegant way and (2) provides insights into the emergence and stability of collective actors.

We then have taken a first step toward a clearer understanding of the limits and applications of cooperative team reasoning by implementing it in a simulation framework that contrasts the agents of orthodox game theory with team reasoners. By simulating environments that are gradually more or less amenable to cooperation, we shed light on the relative performance of team reasoners and show how and under what conditions they outperform classical agents in a range of situations. However, this is only true as long as classical reasoners are not enabled to learn about their game partners. As soon as the individual reasoners come into play, the advantage of team reasoning disappears. Note, however, that the higher the proportion of team reasoners in cooperation settings, the higher the total payoff. In this perspective, team reasoning may not necessarily be beneficial for the individual team reasoner but commonly so for populations that contain groups of team reasoners in them.

References

Ade, Leyla, and Olivier Roy. 2023. Team Reasoning from an Evolutionary Perspective: Categorization and Fitness. *Historical Social Research* 48 (3): 75-93. doi: 10.12759/hsr.48.2023.26.

Amadae, Sonja M., and Daniel Lempert. 2015. The Long-Term Viability of Team Reasoning. *Journal of Economic Methodology* 22 (4): 462-78. doi: 10.1080/1350178X.2015.1024880.

Andreoni, James, and John H. Miller. 1993. Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence. *The Economic Journal* 103 (418): 570-85. doi: 10.2307/2234532.

Angus, Simon D., and Jonathan Newton. 2015. Emergence of Shared Intentionality Is Coupled to the Advance of Cumulative Culture. Edited by Stefano Allesina. *PLOS Computational Biology* 11 (10): e1004587. doi: 10.1371/journal.pcbi.1004587.

Bacharach, Michael. 1999. Interactive Team Reasoning: A Contribution to the Theory of Co-Operation. *Research in Economics* 53 (2): 117-47. doi: 10.1006/reec.1999.0188.

Bacharach, Michael. 2006. *Beyond Individual Choice: Teams and Frames in Game Theory*, ed. Natalie Gold and Robert Sugden. Princeton, N.J: Princeton University Press.

Bacharach, Michael, and Dale O. Stahl. 2000. Variable-Frame Level-n Theory. *Games and Economic Behavior* 32 (2): 220-46. doi:10.1006/game.2000.0796.

Borji, Ali, and Laurent Itti. 2013. Bayesian Optimization Explains Human Active Search. In *Advances in Neural Information Processing Systems*. Vol. 26. Curran Associates, Inc. https://papers.nips.cc/paper_files/paper/2013/hash/a3f390d88e4c41f2747bfa2f1b5f87db-Abstract.html (Accessed May 3, 2023).

Camerer, Colin. 2003. *Behavioral Game Theory: Experiments in Strategic Interaction. The Roundtable Series in Behavioral Economics.* New York, N.Y.; Princeton, N.J: Russell Sage Foundation; Princeton University Press.

Coleman, James S. 1974. *Power and the Structure of Society.* New York: Norton.

Colman, Andrew M., Briony D. Pulford, and Catherine L. Lawrence. 2014. Explaining Strategic Coordination: Cognitive Hierarchy Theory, Strong Stackelberg Reasoning, and Team Reasoning. *Decision* 1 (1): 35-58. doi: 10.1037/dec0000001.

Elsenbroich, Corinna, and Nicolas Payette. 2020. Choosing to Cooperate: Modelling Public Goods Games with Team Reasoning. *Journal of Choice Modelling* 34 (March): 100203. doi: 10.1016/j.jocm.2020.100203.

Gehring, Thomas. 2023. International Organizations as Group Actors. How Institutional Procedures Create Organizational Independence without Delegation to Institutional Agents. *Historical Social Research* 48 (3): 94-124. doi: 10.12759/hsr.48.2023.27.

Gehring, Thomas, and Johannes Marx. 2023. The Emergence and Effects of Non-hierarchical Collective Agency. *Historical Social Research* 48 (3): 7-39. doi: 10.12759/hsr.48.2023.23.

Gilbert, Margaret. 2023. Real Team Reasoning, *Historical Social Research* 48 (3): 40-55. doi: 10.12759/hsr.48.2023.24.

Gold, Natalie. 2017. Team Reasoning: Controversies and Open Research Questions. In *The Routledge Handbook of Collective Intentionality*, ed. Marija Jankovic and Kirk Ludwig, 1st ed., 221-32. Routledge. doi: 10.4324/9781315768571-21.

Gold, Natalie, and Robert Sugden. 2007. Collective Intentions and Team Agency. *The Journal of Philosophy* 104 (3): 109-37.

Hagberg, Aric, Pieter Swart, and Daniel Chult. 2008. Exploring Network Structure, Dynamics, and Function Using NetworkX. United States. https://www.osti.gov/servlets/purl/960616 (Accessed February 14, 2023).

Heuer, Leonie, and Andreas Orland. 2019. Cooperation in the Prisoner's Dilemma: An Experimental Comparison between Pure and Mixed Strategies. *Royal Society Open Science* 6 (7). Royal Society: 182142. doi: 10.1098/rsos.182142.
Karpus, Jurgis, and Natalie Gold. 2016. Team Reasoning: Theory and Evidence. In *The Routledge Handbook of Philosophy of the Social Mind*. Milton Park, Abingdon, Oxon: Routledge.
Knight, Vincent, and James Campbell. 2018. Nashpy: A Python Library for the Computation of Nash Equilibria. *Journal of Open Source Software* 3 (30): 904. doi: 10.21105/joss.00904.
Lewis-Martin, Jimmy. 2022. What Kinds of Groups Are Group Agents? *Synthese* 200 (4): 283. doi: 10.1007/s11229-022-03766-z.
McKelvey, Richard D., Andrew M. McLennan, and Theodore L. Turocy. 2006. Gambit: Software Tools for Game Theory. Version 0.2006. 01.20.
Meincke, Anne Sophie. 2019. The Disappearance of Change: Towards a Process Account of Persistence. *International Journal of Philosophical Studies* 27 (1): 12-30. doi: 10.1080/09672559.2018.1548634.
Newton, Jonathan. 2017. Shared Intentions: The Evolution of Collaboration. *Games and Economic Behavior* 104 (July): 517-34. doi: 10.1016/j.geb.2017.06.001.
Nyikosa, Favour M., Michael A. Osborne, and Stephen J. Roberts. 2018. *Bayesian Optimization for Dynamic Problems*. arXiv. http://arxiv.org/abs/1803.03432 (Accessed July 5, 2022).
Paolo, Ezequiel Di, and Evan Thompson. 2014. The Enactive Approach. In *The Routledge Handbook of Embodied Cognition*. Milton Park, Abingdon, Oxon: Routledge.
Paternotte, Cédric. 2018. The Efficiency of Team Reasoning. *Revue d'économie politique* 128 (3): 447-68. doi: 10.3917/redp.283.0447.
Radzvilas, Mantas, and Jurgis Karpus. 2021. Team Reasoning without a Hive Mind. *Research in Economics* 75 (4): 345-53. doi: 10.1016/j.rie.2021.09.003.
Snoek, Jasper, Hugo Larochelle, and Ryan P Adams. 2012. Practical Bayesian Optimization of Machine Learning Algorithms. In *Advances in Neural Information Processing Systems*, 2951-2959. Vol. 25. Curran Associates, Inc. https://proceedings.neurips.cc/paper/2012/hash/05311655a15b75fab86956663e1819cd-Abstract.html (Accessed April 3, 2022).
Sugden, Robert. 2000. Team Preferences. *Economics & Philosophy* 16 (2): 175-204. doi: 10.1017/S0266267100000213.
Sugden, Robert. 2015. Team Reasoning and Intentional Cooperation for Mutual Benefit. *Journal of Social Ontology* 1 (1): 143-66. doi: 10.1515/jso-2014-0006.
Vanberg, Viktor. 1982. *Markt und Organisation. Individualistische Sozialtheorie und das Problem korporativen Handelns*. Tübingen: Mohr Siebeck.

HSR
Historical Social Research

All articles published in HSR Special Issue 48 (2023) 3:
The Emergence and Effects of Non-hierarchical Collective Agency

HSR
Historical Social Research

All articles published in HSR Special Issue 48 (2023) 3:
The Emergence and Effects of Non-hierarchical Collective Agency