

Deconstructing Cooperation and Ostracism via Multi-Agent Reinforcement Learning

Atsushi Ueshima
Tohoku University
Sendai, Japan
ueshima73@gmail.com

Shayegan Omidshafiei
Google Research
Cambridge, MA, USA
omidshafiei@google.com

Hirokazu Shirado
Carnegie Mellon University
Pittsburgh, PA, USA
shirado@cmu.edu

ABSTRACT

Cooperation is challenging in biological systems, human societies, and multi-agent systems in general. While a group can benefit when everyone cooperates, it is tempting for each agent to act selfishly instead. Prior human studies show that people can overcome such social dilemmas while choosing interaction partners, i.e., strategic network rewiring. However, little is known about how agents, including humans, can learn about cooperation from strategic rewiring and vice versa. Here, we perform multi-agent reinforcement learning simulations in which two agents play the Prisoner’s Dilemma game iteratively. Each agent has two policies: one controls whether to cooperate or defect; the other controls whether to rewire connections with another agent. This setting enables us to disentangle complex causal dynamics between cooperation and network rewiring. We find that network rewiring facilitates mutual cooperation even when one agent always offers cooperation, which is vulnerable to free-riding. We then confirm that the network-rewiring effect is exerted through agents’ learning of ostracism, that is, connecting to cooperators and disconnecting from defectors. However, we also find that ostracism alone is not sufficient to make cooperation emerge. Instead, ostracism emerges from the learning of cooperation, and existing cooperation is subsequently *reinforced* due to the presence of ostracism. Our findings provide insights into the conditions and mechanisms necessary for the emergence of cooperation with network rewiring.

KEYWORDS

Cooperation, Ostracism, Network rewiring, Multi-agent reinforcement learning

1 INTRODUCTION

Understanding how cooperation arises and is sustained is crucial to comprehending biological systems, human societies, and multi-agent systems in general [2, 32, 34]. Cooperation can be difficult to achieve when it creates a social dilemma [5]. In such situations, groups can benefit when everyone cooperates, but it is tempting for each agent to act selfishly instead. Hence, encouraging individuals to develop and sustain cooperation presents significant challenges [11, 21, 36]. How can individual agents learn to overcome short-sighted selfishness and achieve mutual cooperation for the greater good?

One possible way to address social dilemmas is network rewiring [1, 22, 37, 38]. Individuals might not necessarily be bound to interact with their counterparts every time; instead, they might have some agency to establish and sever connections and choose whether to interact with them next time. Network rewiring offers an additional means of responding to the past actions of others. Not only can

agents reciprocate by strategically changing their own cooperation behavior, but they can also alter their network connections, engaging in "tie reciprocity" or "ostracism" [8, 14, 45]. This paper defines ostracism as a specific network-rewiring policy in which agents connect to cooperators and disconnect from defectors. Prior simulations and human-subject experiments demonstrate that cooperation emerges with network rewiring, and most individuals exhibit such strategic network rewiring during the process [7, 10, 28, 29, 38, 45].

However, little is known about how agents spontaneously learn about cooperative policies from network rewiring and vice versa. Previous methodology has limitations in disentangling complex causal dynamics between cooperation and network rewiring. Human-subject experiments, for instance, do not give participants sufficient time and iterations to develop policies through interaction with others [24, 38, 45]. Similarly, computational simulations often use a predetermined set of behavioral policies for agents in cooperation and network-rewiring decisions [13, 31, 32, 37], constraining their ability to examine policy generation. Recent works using multi-agent reinforcement learning (MARL) allow agents to generate behavioral policies in cooperative domains [6, 20, 21, 27]. However, they seek to design interactions and optimization mechanisms at the group level to explicitly facilitate cooperation, rather than investigate *why* cooperation emerges in certain conditions.

Here, we investigate how cooperation and ostracism are learned through the Iterated Prisoner’s Dilemma (PD) [2] in a two-player MARL environment. Two agents play the game using two policies: one policy controls whether to cooperate or defect with the other agent in a PD payoff structure (i.e., *interaction policy*), while the other policy controls whether to rewire (establish or sever) connections with the other agent (i.e., *network-rewiring policy*). Each agent has two distinct neural networks for each policy that are updated through dyadic interactions. In contrast to prior works, extrinsic institutions favoring cooperation are not present in our setting. We manipulate two key dimensions in this study: 1) the frequency of network-rewiring opportunities and 2) the interaction or network-rewiring policy for one agent. By locking one agent’s policy, we examine how network rewiring encourages or discourages agents from learning about cooperation and ostracism from their counterparts. We evaluate the cooperation level of the two agents across the treatments as a consequence of reinforcement learning.

Our findings show that network rewiring facilitates mutual cooperation even when one agent always offers cooperation, which is vulnerable to free-riding [32]. We also confirm that the network-rewiring effect is exerted through agents’ learning of ostracism. However, ostracism alone is not sufficient to make cooperation

emerge. Instead, ostracism emerges from the learning of cooperation, and existing cooperation is subsequently *reinforced* due to the presence of ostracism. Our study provides valuable insights into the conditions and mechanisms necessary for the emergence of cooperation with network rewiring.

Our work makes the following contributions: 1) we develop a MARL framework to explore the coevolution of cooperation and population structure, 2) we uncover the complex cause-effect relationships involved in the dual learning of cooperation and network rewiring, and 3) we demonstrate the potential application of MARL in addressing important questions in system biology, social sciences, and humanities.

2 RELATED WORK

Evolutionary games. Evolutionary game theory has been used to analyze cooperation difficulties by extending traditional game theory to long-term, multi-player scenarios [32, 46]. For instance, Nowak and May utilized the theoretical framework to demonstrate that the evolution of behavioral policies alone is insufficient to alter social dynamics that attract individuals to defection [31]. Their research and subsequent studies suggest that certain interaction structures, such as network topology [33] and heterogeneity [41], could help cooperators spread over an entire population while resisting defection. Several other works have applied coevolution rules to evolutionary games, indicating that cooperation can emerge as interaction networks evolve [10, 37].

Evolutionary games originally focus on intergenerational evolution, where agents with higher fitness reproduce their strategies more frequently in a population [32]. These simulations prepare for two or more strategies (or fixed policies) to interact with each other and update the fraction of each strategy based on their earnings or fitness. However, this assumption and setting do not allow us to scrutinize the emergence of cooperation in the real world through social learning. In real-world interindividual interactions, agents cannot directly observe or *exactly* mimic others’ policies. We bypass this issue in our study by applying MARL to network-rewiring settings, allowing agents to learn both policies to cooperate and to change interaction connections by observing an environment (including the other agent’s behaviors) and obtaining rewards.

Human studies. Evolutionary game theory has been tested through human-subject experiments to study human cooperation [39]. By allowing human subjects to play a model game accordingly, researchers have examined what can make people cooperate. For instance, Rand, Arbesman, and Christakis conducted a lab experiment where human participants played a multi-player Prisoner’s Dilemma game repeatedly, changing their interaction partners in a sizable group [38]. The results showed that cooperation stabilized at a high level through network rewiring while it decayed over time when their partners were shuffled or fixed. Shirado et al. clarified that people establish connections with cooperators and sever them from defectors, and the specific network-rewiring policy helps cooperators cluster and benefit each other [45].

Human-subject experiments have also revealed population-wise differences in cooperation [12], wherein some populations admit egocentrism, while others show altruism in cooperation games. These findings suggest that behavioral experiments allow us to

study what policies humans have, but not how humans develop such policies through interaction. The experimental sessions, which usually last for 30-45 minutes (and for 20 days at longest [24]), are simply not enough to examine the longer-term processes of developing and establishing individual policies.

Multi-agent reinforcement learning. MARL has been used to overcome the limitations of traditional approaches toward understanding the emergence of cooperation [3, 4, 9, 16, 18, 19, 21, 23, 27, 35, 40]. For instance, McKeen et al. have employed MARL to explore network-rewiring strategies for intervention agents and discovered more effective intervention policies than simple ones [27]. Including McKeen’s work, however, most studies incorporating MARL use hyper-parameters or reward functions to incentivize agents towards inducing cooperation, thereby not exploring different scenarios under which cooperation does or does not arise.

On the other hand, Lupu and Precup have extended the action space of agents, instead of employing group-level optimization, to identify the conditions that facilitate the learning of cooperation [23]. They incorporated a peer rewarding mechanism in their MARL simulations, which significantly improved learning progression for cooperation. Our work further extends their approach to make agents learn peer rewarding and punishing policies through network rewiring, such as attachment and detachment, while simultaneously learning cooperation.

3 METHODS

We conduct our studies in a MARL environment wherein a pair of agents simulate network-rewiring interactions across multiple episodes (Figure 1). Each agent decides their behavior based on their observations, intending to maximize their respective rewards. To examine how network rewiring affects the spontaneous development of cooperation, we do not incorporate any incentive mechanisms that maximize group benefits. We perform the simulations using Acme, a framework for distributed reinforcement learning [15]. We next formulate our approach in detail.¹

3.1 Formulation

3.1.1 Policies and action space. Each agent in our framework makes use of two policies to make decisions at each timestep: an *interaction policy* and a *network-rewiring policy* (see Figure 1).

Network-rewiring policy: at each timestep of a given episode, agents may have an opportunity to modify their connections to one another. Given such a network-rewiring opportunity, each agent can choose to establish or sever connections with the other agent. We conduct ablations over the number of network-rewiring opportunities in a given episode (e.g., agents in some experiments are allowed to rewire at every timestep, and in others only allowed to rewire at intermittent timesteps; see Section 3.2.1 for details). When both agents decide to establish a connection at the same timestep, the connection is established (or maintained if already present; i.e., bilateral tie-making). Otherwise, the connection is not established (or removed if already present; i.e., unilateral tie-breaking).

¹We will open source our codebase upon acceptance.

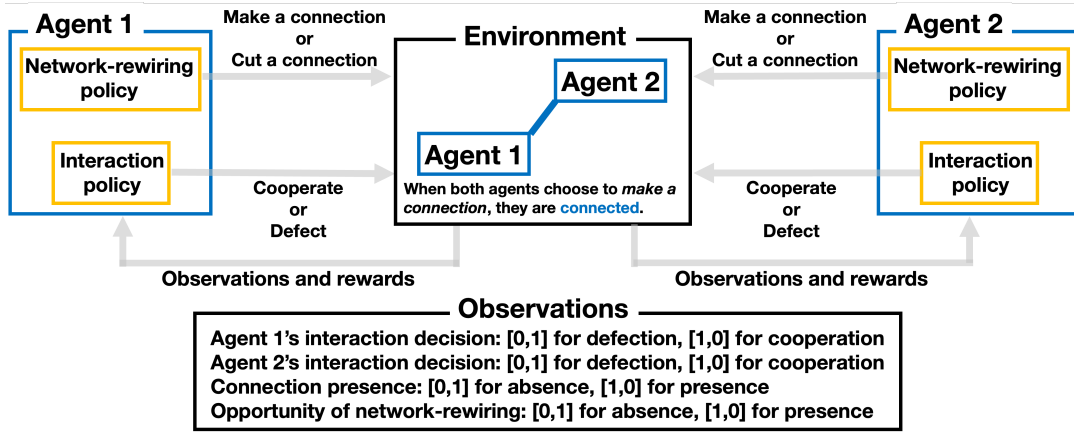


Figure 1: The overview of our simulation setting. Each agent has two sets of Q-networks: one for the interaction policy and the other for the network-rewiring policy. In the environment, two agents are connected when both choose to make a connection in a given timestep (i.e., bilateral tie-making), and are otherwise disconnected (i.e., unilateral tie-breaking). Agents play a round of Prisoner’s Dilemma in a given timestep if they are connected and receive payoffs accordingly (and otherwise receive a payoff of 0 if no connection exists). Network-rewiring policy first affects the environment by making or breaking a connection before the interaction policy takes effect, thus making it possible for one agent to prevent even the possibility of cooperation at each timestep.

Table 1: Payoff structure of the Prisoner’s Dilemma game. Each cell indicates Agent 1’s payoff (left) and Agent 2’s payoff (right) with indicated interaction options.

		Agent 2	
		Cooperation	Defection
Agent 1	Cooperation	1, 1	-1, 2
	Defection	2, -1	0, 0

Interaction policy: depending on the connection status of the agents, they may then engage in a game with one another. If agents are disconnected, no such interaction occurs between them. If agents are connected, then for that timestep they can each choose to cooperate or defect in a Prisoner’s Dilemma scenario as described in Table 1. This interaction structure creates a social dilemma that favors defection over cooperation: if agents simply learn a short-sighted, greedy policy, the dyad archives the equilibrium of mutual defection that gives both no payoffs. It is, however, a worse outcome than rewarding mutual cooperation.

Moreover, to earn the cooperation benefit, it is necessary for agents to be connected. By establishing connections, however, agents are at risk of being exploited by defectors. Thus, the challenge with this interaction setting is how agents can learn to overcome myopia and exploitation risks so as to obtain public goods through network connections and mutual cooperation.

3.1.2 *Observation space.* At each timestep, each agent receives the following set of one-hot observations:

- The agent’s own previous interaction decision ([0,1] for defection, [1,0] for cooperation).

- The other agent’s previous interaction decision ([0,1] for defection, [1,0] for cooperation).
- Whether a network edge was present in the previous timestep: ([0,1] for absence, [1,0] for presence).
- Whether there was an opportunity to perform network rewiring in the previous timestep ([0,1] for absence, [1,0] for presence).

For the last observation, agents always receive [0,1] in the first timestep of each episode because they do not have the previous timestep or the opportunity to perform network rewiring. See also the Supplementary Material for details regarding the observational space in the first timestep of each episode.

3.1.3 *Episode sequence.* For each experiment trial, agents interact in 200,000 episodes, each consisting of 10 timesteps (i.e., 2,000,000 steps total). Each episode starts with two agents connected. In each timestep, the two policy networks of each agent provide network-rewiring and interaction actions, respectively (Figure 1). If agents are given a network-rewiring opportunity (see Section 3.2.1), they first take the network-rewiring actions, after which they are either connected with or disconnected from one another. Otherwise, the connection is the same as the previous timestep. If the agents are connected, they then interact with their interaction actions (i.e., cooperation or defection) and receive payoffs based on their own and others’ decisions (Table 1). However, if they are disconnected, interaction does not occur, and agents receive 0 payoff for that timestep. The process is repeated until the end of the episode.

3.1.4 *Learning algorithm.* For training, agents use the Double Deep Q-Network (DDQN) algorithm [47] along with prioritized experience replay [42]. Each agent has two sets of Q-networks: one for the interaction policy and the other for the network-rewiring policy. The agents’ policies train independently from one another, and

no parameters are shared between the agents. Due to this setting, agents cannot directly copy the other’s policy, even if the other agent performs better. For each Q-network, we use a Multi-Layer Perceptron (MLP) neural network architecture with 16, 16, 2 hidden nodes in each respective layer, and *tanh* activation function. See the Supplementary Material for more information about hyperparameters.

3.2 Treatment ablations

We manipulate two dimensions in this study: 1) the frequency of network-rewiring opportunities and 2) the interaction policy for one of the agents. In some of our experiments, by locking one agent’s policy, we examine how network rewiring encourages or discourages agents from learning about cooperation and ostracism from their counterparts.

In total, we examine 9 treatment combinations of network-rewiring frequency and policy fixation: 3 network-rewiring opportunity conditions (no-rewiring, half-rewiring, and full-rewiring; see Section 3.2.1) cross with 3 fixed policy conditions (no-bias, ALLC-bias, and TFT-bias; see Section 3.2.2). We also explore the ostracism-bias condition with the three levels of network-rewiring frequency. We conduct 15 random seeds for each combinational condition.

3.2.1 Network-rewiring opportunities. We manipulate the frequency of network-rewiring opportunities as follows. In the "no-rewiring" condition, the agents do not have the opportunity to make network-rewiring decisions. Since they are always connected, any learning driven by network rewiring does not occur in this control condition. In the "full-rewiring" condition, on the other hand, the agents have the opportunity to make network-rewiring decisions every timestep. In addition, we explore the "half-rewiring" condition where the agents make network-rewiring decisions every even timestep (i.e., 2, 4, 6, 8, and 10th timesteps) while they make interaction decisions every timestep. Thus, the agents are expected to consider a longer-term consequence of network-rewiring decisions in the half-rewiring condition.

3.2.2 Fixed policies. Independent of the network-rewiring opportunities, we also conduct ablations controlling the policy of one of the agents for *interaction* actions. In the "no-bias" condition, no fixed policy is implemented (i.e., both agents learn their interaction policy). In the "ALLC-bias" condition, one agent has a fixed policy of always cooperating (ALLC) with the other agent. In the "TFT-bias" condition, one agent is forced to follow the Tit-for-Tat (TFT) policy, which starts with cooperation and then copies the other agent’s previous interaction action. In all the conditions, both agents (including the one whose interaction policy is fixed) learn the network-rewiring policy using RL.

We design the fixed policy treatment, following the system-biology approach to explore evolutionary stable strategies [31, 32]. With policy fixation, we examine how cooperation would evolve from the seeds of a cooperative policy and how network rewiring would affect the process. It is widely accepted that ALLC agents are vulnerable to exploitation by defectors and that their cooperative policy rarely generates mutual cooperation in a standard social dilemma setting. On the other hand, TFT agents, although more

complicated, could trigger the development of mutual cooperation when their counterparts seek cooperation [2]. TFT agents are also robust to defections as they prevent themselves from being exploited by copying defection.

In addition, we examine a condition in which we control the policy of one of the agents for *network-rewiring* actions. In the "ostracism-bias" condition, one agent chooses to establish connections when the other cooperates and to serve connections when the other defected in the previous timestep. This additional condition enables us to extend our main findings by examining whether ostracism alone can develop cooperation through RL.

3.3 Evaluation

For each of the treatment ablations above, we evaluate how well agents cooperate with each other. To measure the level of mutual cooperation, we calculate the average number of times the two agents chose to cooperate simultaneously per episode, based on the interaction structure shown in Table 1. We only count mutual cooperation actions when the agents have a connection at a given time. For instance, if both agents are connected and cooperate five times in a 10-timestep episode, the rate of mutual cooperation for that episode is 0.5. See the Supplementary Information for further details.

4 RESULTS

This section presents the results of our experiments, including treatment ablations.

4.1 Effects of network rewiring on cooperation

Figure 2 shows average mutual cooperation rates across the level of network-rewiring capability and fixed interaction policies throughout training. The comparison across the fixed policies shows that cooperation emerges with network rewiring, especially when one agent adopts a cooperative policy, such as ALLC and TFT. In keeping with prior work [37, 38, 45], cooperation is unlikely to arise or persist without network rewiring, except when one agent implements the TFT policy. Moreover, even with network rewiring, cooperation does not emerge among agents who learn policies from scratch (i.e., the no-bias condition) in our simulations. This result highlights the challenge of resolving social dilemmas, especially without at least some individuals who are willing to cooperate.

We next analyze the ALLC-bias column of Figure 2. We find that enabling rewiring opportunities leads to a more frequent emergence of mutual cooperation, especially when one agent consistently chooses cooperation, i.e., the ALLC policy. In the ALLC-bias condition, cooperation collapses in the no-rewiring and half-rewiring conditions due to the other (unlocked-policy) agent taking advantage of the cooperation benefit. Simply having agents implement a generous, cooperative policy is not enough to overcome social dilemmas as free riders easily exploit it. However, if agents are given sufficient chances to rewire their connections, cooperation becomes more resilient to defection (as shown at the "ALLC bias" crossed "Full rewiring" in Figure 2). Even in the full-rewiring condition, cooperation levels actually decrease after around 100K episodes with the ALLC-bias agents, but then bounce back. This result suggests

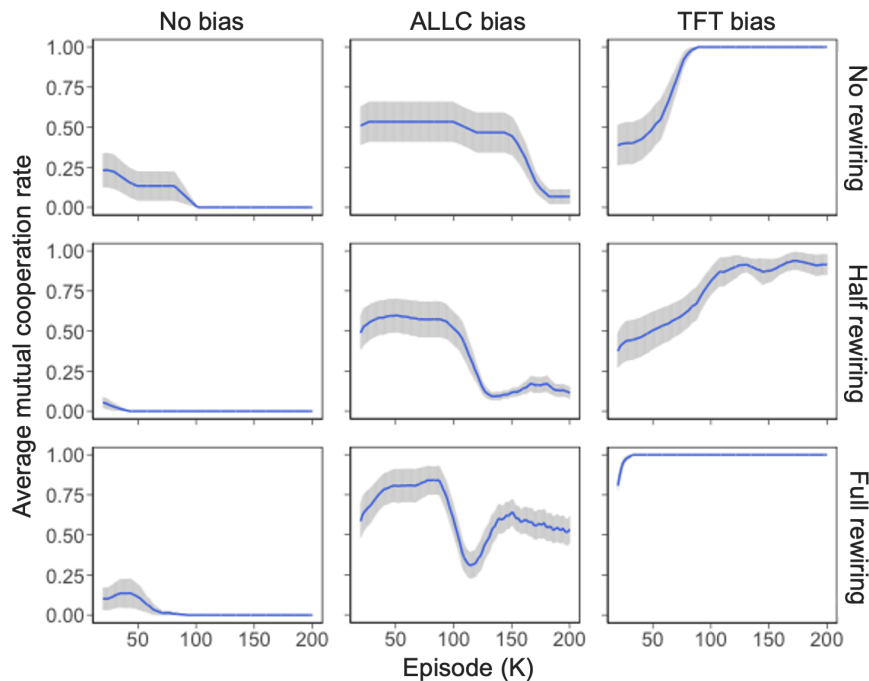


Figure 2: Average mutual cooperation rates across behavioral biases and the opportunity of network rewiring. Blue lines indicate the average rates of mutual cooperation. Gray shades indicate standard errors ($N=15$). The result shows the opportunity of network rewiring facilitates the emergence of mutual cooperation when one of the agents in a group is equipped with any of the behavioral biases, and the ALLC-bias condition was more heavily affected by the amount of rewiring opportunities.

that network rewiring enhances cooperation resilience through reinforcement learning, as cooperators can defend themselves against defectors by detachment.

We next discuss the TFT-bias column of Figure 2. When an agent implements the TFT policy, which involves starting with cooperation and following their counterpart’s interaction, most random seeds result in mutual cooperation across the rewiring conditions. This suggests that mutual cooperation relies less on learning the network-rewiring policy in the TFT-bias condition than in the ALLC-bias condition. The TFT policy allows the agents to prevent exploitation from free riders by copying defection, which means that further opportunities to detach from defectors may be less impactful for the group dynamics.

Nevertheless, network rewiring can accelerate the process of mutual cooperation in conjunction with the TFT policy. In the TFT-bias full-rewiring condition, agents’ interaction actions quickly converge to mutual cooperation and never collapse afterward. It is worth noting that the development of mutual cooperation is slower in the half-rewiring condition than in the no-rewiring condition, possibly due to fewer learning opportunities for their interaction policies. If agents initially sever connections, they earn zero rewards regardless of their interaction actions (i.e., cooperation and defection), which suppresses their learning about how they should interact with each other. Network rewiring affects the developmental process of cooperation by balancing reciprocal opportunities for connections and learning opportunities for cooperation.

4.2 Learned network-rewiring policies

We then turn to how agents develop their network-rewiring policies through interactions with the social dilemma. Figure 3 shows the fractions of connecting to the counterparts (instead of disconnection from them) based on the data obtained from the end phase of training. These results indicate, after substantial interactions and learning, whether an agent establishes or seizes the connection with their counterpart based on their interaction action. Note that timesteps where agents do not have the opportunity to rewire in the half-rewiring condition are excluded from the analysis. For instance, the top-left cell illustrates the results of the no-bias half-rewiring condition. In this condition, agents attempt to connect or maintain a connection with their counterparts 84.6% of the time (instead of disconnecting from them) when their counterparts have cooperated in the last timestep and 85.3% of the time when they have defected.

Our analysis shows that agents generally prefer to establish connections rather than sever them. Regardless of whether their counterparts cooperate or defect, agents in the no-bias and TFT-bias conditions choose to have connections with their counterparts more than 80% of the time. This is an interesting outcome, indicating that agents prefer to connect with each other to enable interactions and the possibility of earning payoffs.

However, the network-rewiring policy correlates with cooperation in the opposite way between the no-bias and TFT-bias conditions. Choosing to connect associates with mutual defection in the




Figure 3: Fraction of agents establishing and severing connections across the conditions after reinforcement learning of cooperation and network rewiring. Blue bars indicate the fractions that agents *intend* to connect to their counterparts, instead of disconnecting from them (red bars). When both agents choose to connect at the same timestep, the connection is established or maintained. Otherwise, the connection is not established or removed. ALLC-biased agents exhibit higher rates of ostracism: attempting to establish connections with the other agent that cooperated and severing connections with the other that defected in the last timestep.

no-bias condition, where most agents start with defection (Figure 2). Defectors face no risk of exploitation, so they seek out interactions for possible earnings, which makes them prefer connecting. It makes it more difficult for agents to change their behavior to cooperation. In contrast, in the TFT-bias condition, agents with the TFT policy can manage the risk of exploitation by choosing defection when facing defectors. This increases the expected benefit of connection even when their counterparts have chosen defection in the last timestep. As mutual cooperation evolves with the TFT policy (Figure 2), the advantage of having connections increases.

In the ALLC-bias scenario, agents develop a more complex policy for network rewiring (Figure 3). Agents with always cooperation learn to ostracize other agents in the full-rewiring condition. They

connect with cooperators 63.4% of the time, while they connect with defectors only 39.9% of the time. This shows that agents eventually value the quality of connections over the quantity during network rewiring. In fact, the learned rewiring bias is similar to how humans behave in similar rewiring opportunities in a cooperation scenario [7, 45].

On the other hand, when agents have half the network-rewiring opportunities, they do not learn to connect with cooperators while they learn to disconnect from defectors. They connect with cooperators only 32.7% of the time, while they connect with defectors 28.5% of the time in the half-rewiring condition. This finding suggests that agents need sufficient rewiring opportunities *vis-a-vis* interaction ones to develop selective rewiring policies based on the other's past actions.

Finally, the ALLC-bias condition does not allow unbiased agents to develop a network-rewiring policy with defectors because their counterparts always choose cooperation, regardless of network-rewiring opportunities.

Our findings suggest that, given sufficient network-rewiring opportunities, unconditional cooperators can learn the policy of ostracism, which leads to resilient cooperation (as shown at the "ALLC bias" crossed "Full rewiring" in Figure 2). In contrast, conditional cooperators, such as TFT agents, rarely learn ostracism (as shown in the "TFT bias" column of Figure 3) because the benefit of connection outweighs its cost. The stable learning of ostracism depends on the unconditional cooperation of agents. Once ostracism is established, it prevents cooperation from dissolution.

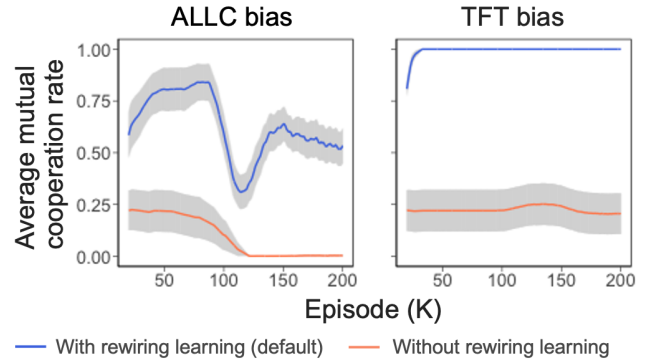


Figure 4: Average mutual cooperation rates in the ALLC-bias and the TFT-bias when the learning of rewiring action is disabled (red lines) and enabled (blue lines). All are in the full-rewiring condition. Lines indicate the average rates. Shades indicate standard errors ($N=15$). We show the results "with rewiring learning" for comparison, which are identical to those in the full-rewiring condition shown in Figure 2.

4.3 The learning of ostracism, rather than the opportunity of rewiring, facilitates cooperation

Next, we deconstruct and pinpoint the causal effects of network rewiring on cooperation. To do so, we disable the learning of the

network-rewiring policy, isolating the effects of rewiring opportunities *per se*. We set the learning rate of the network-rewiring policy to zero for both agents in each group, allowing the agents to exhibit network-rewiring behavior as in the initial random policy. We only permit the agents to update the interaction policy governing whether to cooperate.

As shown in Figure 4, when the updates of the neural network for the network-rewiring policy are disabled, the cooperation level is significantly lower in both the ALLC-bias and TFT-bias conditions compared to when the learning is allowed (which is identical to the result of the full-rewiring condition in Figure 2). Agents develop almost no mutual cooperation without rewiring-policy learning, which is more defective than the interactions without rewiring opportunities (i.e., the no-rewiring condition).

This is because, without learning, agents randomly choose to connect and disconnect, and consequently, they mutually establish connections and interact only 25% of the time on average. Thus, the maximum level of mutual cooperation is limited to 25%. Furthermore, network-rewiring opportunities solely reduce the opportunities to learn their interaction policies by intermittent disconnections. This is salient in both bias conditions. In the ALLC-bias condition, agents resume mutual cooperation after a halfway collapse with network-rewiring learning, but they do not without learning. In the TFT-bias condition, the level of mutual cooperation rises to 100% in the first 30K episodes with network-rewiring learning, while no such increase occurs without learning.

The development and sustainability of cooperation through network rewiring is due to agents developing a specific network-rewiring policy, such as ostracism, through learning from interactions. In other words, if agents lack the ability to learn an effective way to rewire connections, they will not achieve mutual cooperation even with network rewiring. Thus, the learning of ostracism, rather than the opportunity for rewiring, matters in fostering cooperation.

4.4 Cooperation does not emerge from ostracism alone

Our findings suggest that learning network-rewiring policies like ostracism is critical in addressing social dilemmas. Acquiring the ability to cooperate first as a behavioral predisposition facilitates the acquisition of ostracism behaviors, leading to the emergence of mutual cooperation. However, it remains unclear whether ostracism can foster cooperation in the absence of conditional or unconditional cooperators. To address this, we conduct additional simulations in the "ostracism-bias" condition, where one agent always aims to connect with the other that cooperated in the previous timestep, and disconnects from those that defected earlier.

We find that in the ostracism-bias condition, agents did not establish mutual cooperation in any of the rewiring conditions (Figure 5). The agents with ostracism abilities can defect just like the no-bias agents. Thus, as in the no-bias condition, the interaction quickly leads to mutual defection in the ostracism-bias condition, providing no opportunity for agents to experience selective network rewiring through ostracism. These findings suggest that *extrinsically* endowing agents with the capacity to ostracize others does not enable mutual cooperation. Instead, as shown in the previous section, it

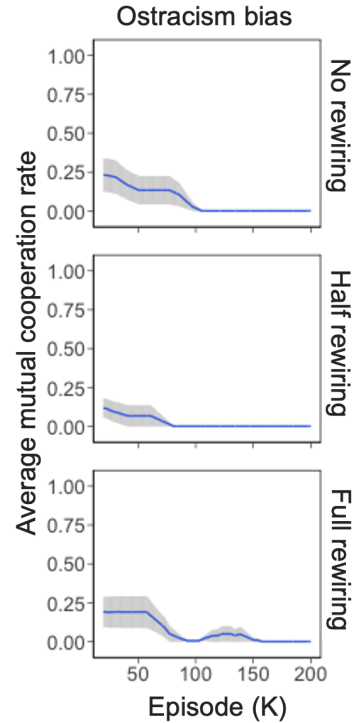


Figure 5: Average mutual cooperation rates in the ostracism-bias condition across the rewiring conditions. In the ostracism-bias condition, one of two agents in a group always chooses to establish connections when the other cooperates and to sever connections when the other defected in the previous timestep. This result suggests that cooperation rarely arises from ostracism alone.

is crucial for agents to learn ostracism behaviors *spontaneously* through cooperative bias to develop and sustain mutual cooperation.

5 DISCUSSION

Our study shows that the ability to rewire connections for interaction is highly correlated with the emergence of mutual cooperation. When agents have more opportunities to establish and sever their connections, they achieve mutual cooperation more often. This finding is in keeping with prior human studies showing that cooperation emerges with network rewiring [27, 28, 38, 45]. By using a two-policy approach and MARL, our study further clarifies that the network-rewiring effect is exerted in the presence of behavioral biases [31, 32]. In our simulations, only when one agent implements a simple cooperative policy (ALLC) or complex reciprocal one (TFT), can agents have mutual cooperation with rewiring capabilities. Building on such cooperative agents, our study shows that network rewiring can make cooperation resilient against defection.

We find that network rewiring has both advantages and disadvantages in the evolution of cooperation. As prior work suggested, enabling network rewiring can help cooperators prevent defectors

from taking advantage of them [38, 45]. On the other hand, our study shows that it can also limit the opportunities for agents to interact with each other, earn payoffs, and learn their policies. For example, network rewiring delays the development of cooperation in the half-rewiring TFT-bias condition, suggesting the disadvantages outweigh the advantages at lower connection mobility. Our findings suggest that for agents to use the benefits of rewiring opportunities for cooperation, they might need to learn a selective rewiring policy, such as ostracism, through limited interaction opportunities. It could also be favorable for cooperation that a dynamic network consists of diverse individuals with biased policies, including unconditional or conditional cooperators [41].

We clarify the role of ostracism in underlying cooperation. In our study, when network-rewiring opportunities are present, cooperation-biased agents are more effective at learning ostracism than agents without a behavioral bias or with the TFT policy. While no-bias and TFT agents prefer connecting to their counterparts, regardless of their interaction actions, always-cooperating (ALLC) agents choose whether to connect based on whether their counterparts cooperate or not. This finding suggests that unconditional cooperators or "zealots" might be necessary to establish ostracism, which facilitates group cooperation [26, 44]. We also confirm that cooperation rarely emerges from ostracism alone. Rather, ostracism can act as a catalyst that helps dynamic networks favor cooperation over defection.

There are other features potentially relevant to the coevolution of cooperation and ostracism that we have not explored in this study. For example, we do not observe mutual cooperation in the no-bias condition, even with the opportunity for network rewiring. However, our findings also suggest that mutual cooperation would be possible if at least one agent learned, for example, the TFT policy. It is an important next step to address how agents can learn such a complex, reciprocal policy through interaction and network rewiring in the setting we considered. We find that agents can learn ostracism, which is a reciprocal policy for network rewiring, from a simple cooperation bias. Thus, the TFT might also be developed from simple behavioral biases.

Additionally, our study only considers symmetric interactions between two agents, as this setting permits a more scrutinous analysis of inter-agent interactions. Nonetheless, expanding the setting to larger multi-agent systems with complex network structures and dynamics could provide further insights into the coevolution of cooperation and ostracism. Moreover, other interaction strictures, such as asymmetric cooperation benefit and cost between agents [17, 30] or explicit costs of network connections [25], might drive reinforcement learning and interaction to different equilibria with network rewiring. Finally, the learning process in our simulations is guided by an off-policy algorithm, the Double Deep Q-Network (DQN). Thus, agents learn their policies partially independently from their latest observations. Future studies could compare the potential differences between off-policy and on-policy algorithms, such as Proximal Policy Optimization (PPO) [43], to further understand how learning processes can influence the emergence of cooperation and ostracism.

In conclusion, our study highlights the importance of network rewiring and the learning of ostracism in the coevolution of cooperation and ostracism in MARL settings. Network-rewiring opportunities could help a few individuals having a cooperative tendency learn ostracism, which facilitates, in turn, cooperation over a population. These findings contribute to understanding how and when different types of policies, such as interaction and network-rewiring, could emerge simultaneously in multi-agent systems with network dynamics.

REFERENCES

- [1] Erol Akçay. 2018. Collapse and rescue of cooperation in evolving dynamic networks. *Nat. Commun.* 9, 1 (July 2018), 1–9.
- [2] R Axelrod and W D Hamilton. 1981. The evolution of cooperation. *Science* 211, 4489 (March 1981), 1390–1396.
- [3] Kris Cao, Angeliki Lazaridou, Marc Lanctot, Joel Z Leibo, Karl Tuyls, and Stephen Clark. 2018. Emergent communication through negotiation. *arXiv preprint arXiv:1804.03980* (2018).
- [4] Caroline Claus and Craig Boutilier. 1998. The Dynamics of Reinforcement Learning in Cooperative Multiagent Systems. In *Proceedings of the Fifteenth National/Tenth Conference on Artificial Intelligence/Innovative Applications of Artificial Intelligence* (Madison, Wisconsin, USA) (AAAI '98/IAAI '98). American Association for Artificial Intelligence, USA, 746–752.
- [5] Robyn M Dawes. 1980. Social Dilemmas. *Annu. Rev. Psychol.* 31, 1 (Jan. 1980), 169–193.
- [6] Enrique Munoz de Cote, Alessandro Lazaric, and Marcello Restelli. 2006. Learning to cooperate in multi-agent social dilemmas. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems* (Hakodate, Japan) (AAMAS '06). Association for Computing Machinery, New York, NY, USA, 783–785.
- [7] Katrin Fehl, Daniel J van der Post, and Dirk Semmann. 2011. Co-evolution of behaviour and social network structure promotes human cooperation. *Ecol. Lett.* 14, 6 (April 2011), 546–551.
- [8] Matthew Feinberg, Robb Willer, and Michael Schultz. 2014. Gossip and Ostracism Promote Cooperation in Groups. *Psychol. Sci.* 25, 3 (Jan. 2014), 656–664.
- [9] Jakob Foerster, Ioannis Alexandros Assael, Nando De Freitas, and Shimon Whiteson. 2016. Learning to communicate with deep multi-agent reinforcement learning. *Advances in neural information processing systems* 29 (2016).
- [10] Feng Fu, Christoph Hauert, Martin A Nowak, and Long Wang. 2008. Reputation-based partner choice promotes cooperation in social networks. *Physical Review E* 78, 2 (Aug. 2008), 026117–026118.
- [11] Garrett Hardin. 1968. The Tragedy of the Commons. *Science* 162, 3859 (1968), 1243–1248.
- [12] Joseph Heinrich, R Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Herbert Gintis, and Richard McElreath. 2001. In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies. *Economic and social behavior* 91, 2 (May 2001), 73–78.
- [13] Christian Hilbe, Martin A Nowak, and Karl Sigmund. 2013. Evolution of extortion in Iterated Prisoner's Dilemma games. *Proc. Natl. Acad. Sci. U. S. A.* 110, 17 (April 2013), 6913–6918.
- [14] David Hirshleifer and Eric Rasmusen. 1989. Cooperation in a repeated prisoners dilemma with ostracism. *Journal of Economic Behavior and Organization* 12 (April 1989), 87–106.
- [15] Matthew W. Hoffman, Bobak Shahriari, John Aslanides, Gabriel Barth-Marón, Nikola Momchev, Danila Sinopalnikov, Piotr Stańczyk, Sabela Ramos, Anton Raichuk, Damien Vincent, Léonard Hussenot, Robert Dadashi, Gabriel Dulac-Arnold, Manu Orsini, Alexis Jacq, Johan Ferret, Nino Vieillard, Seyed Kamyar Seyed Ghasemipour, Sertan Girgin, Olivier Pietquin, Feryal Behbahani, Tamara Norman, Abbas Abdolmaleki, Albin Cassirer, Fan Yang, Kate Baumli, Sarah Henderson, Abe Friesen, Ruba Haroun, Alex Novikov, Sergio Gómez Colmenarejo, Serkan Cabi, Caglar Gulcehre, Tom Le Paine, Srivatsan Srinivasan, Andrew Cowie, Ziyu Wang, Bilal Piot, and Nando de Freitas. 2020. Acme: A Research Framework for Distributed Reinforcement Learning. *arXiv preprint arXiv:2006.00979* (2020). <https://arxiv.org/abs/2006.00979>
- [16] Edward Hughes, Joel Z Leibo, Matthew Phillips, Karl Tuyls, Edgar Dueñez-Guzman, Antonio García Castañeda, Iain Dunning, Tina Zhu, Kevin McKee, Raphael Koster, et al. 2018. Inequity aversion improves cooperation in intertemporal social dilemmas. *Advances in neural information processing systems* 31 (2018).
- [17] Marco A Janssen, John M Anderies, and Sanket R Joshi. 2011. Coordination and cooperation in asymmetric commons dilemmas. *Exp. Econ.* 14, 4 (Nov. 2011), 547–566.
- [18] Spiros Kapetanakis and Daniel Kudenko. 2002. Reinforcement Learning of Coordination in Cooperative Multi-Agent Systems. In *Eighteenth National Conference*

- on *Artificial Intelligence* (Edmonton, Alberta, Canada). American Association for Artificial Intelligence, USA, 326–331.
- [19] Dong-Ki Kim, Miao Liu, Shayegan Omidshafiei, Sebastian Lopez-Cot, Matthew Riemer, Gerald Tesaro, Murray Campbell, Sami Mourad, Golnaz Habibi, and Jonathan P How. 2020. Heterogeneous knowledge transfers via hierarchical teaching in cooperative multiagent reinforcement learning. In *Proc of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*. 1–8.
- [20] Raphael Köster, Dylan Hadfield-Menell, Richard Everett, Laura Weidinger, Gillian K Hadfield, and Joel Z Leibo. 2022. Spurious normativity enhances learning of compliance and enforcement behavior in artificial agents. *Proc. Natl. Acad. Sci. U. S. A.* 119, 3 (Jan. 2022).
- [21] Joel Z Leibo, Vinicius Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. 2017. Multi-agent Reinforcement Learning in Sequential Social Dilemmas. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (São Paulo, Brazil) (AAMAS '17)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 464–473.
- [22] Aming Li, Lei Zhou, Qi Su, Sean P Cornelius, Yang-Yu Liu, Long Wang, and Simon A Levin. 2020. Evolution of cooperation on temporal networks. *Nat. Commun.* 11, 1 (April 2020), 1–9.
- [23] Andrei Lupu and Doina Precup. 2020. Gifting in Multi-Agent Reinforcement Learning. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '20)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 789–797.
- [24] Andrew Mao, Lili Dworkin, Siddharth Suri, and Duncan J Watts. 2017. Resilient cooperators stabilize long-run cooperation in the finitely repeated Prisoner's Dilemma. *Nat. Commun.* 8, 1 (Jan. 2017), 13800.
- [25] Gerald Marwell, Pamela E Oliver, and Ralph Prael. 1988. Social Networks and Collective Action: A Theory of the Critical Mass. III. *Am. J. Sociol.* 94, 3 (Nov. 1988), 502–534.
- [26] Naoki Masuda. 2012. Evolution of cooperation driven by zealots. *Sci. Rep.* 2 (Sept. 2012).
- [27] Kevin R McKee, Andrea Tacchetti, Michiel A Bakker, Jan Balaguer, Lucy Campbell-Gillingham, Richard Everett, and Matthew Botvinick. 2023. Scaffolding cooperation in human groups with deep reinforcement learning. *Nat Hum Behav* (Sept. 2023).
- [28] John M McNamara, Zoltan Barta, Lutz Fromhage, and Alasdair I Houston. 2008. The coevolution of choosiness and cooperation. *Nature* 451, 7175 (Jan. 2008), 189–192.
- [29] David Melamed, Ashley Harrell, and Brent Simpson. 2018. Cooperation, clustering, and assortative mixing in dynamic networks. *Proc. Natl. Acad. Sci. U. S. A.* 115, 5 (Jan. 2018), 951–956.
- [30] Akihiro Nishi, Hirokazu Shirado, David G Rand, and Nicholas A Christakis. 2015. Inequality and visibility of wealth in experimental social networks. *Nature* 526, 7573 (Jan. 2015), 426–429.
- [31] M Nowak and K Sigmund. 1993. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature* 364, 6432 (July 1993), 56–58.
- [32] Martin A Nowak. 2006. *Evolutionary Dynamics*. Harvard University Press.
- [33] Hisashi Ohtsuki, Christoph Hauert, Erez Lieberman, and Martin A Nowak. 2006. A simple rule for the evolution of cooperation on graphs and social networks. *Nature* 441, 7092 (May 2006), 502–505.
- [34] Mancur Olson. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Harvard University Press.
- [35] Shayegan Omidshafiei, Jason Pazis, Christopher Amato, Jonathan P How, and John Vian. 2017. Deep decentralized multi-task multi-agent reinforcement learning under partial observability. In *International Conference on Machine Learning*. PMLR, 2681–2690.
- [36] Elinor Ostrom. 2000. Collective Action and the Evolution of Social Norms. *J. Econ. Perspect.* 14, 3 (Sept. 2000), 137–158.
- [37] Matjaž Perc and Attila Szolnoki. 2010. Coevolutionary games—A mini review. *Biosystems.* 99, 2 (Feb. 2010), 109–125.
- [38] David G Rand, Samuel Arbesman, and Nicholas A Christakis. 2011. Dynamic social networks promote cooperation in experiments with humans. *Proc. Natl. Acad. Sci. U. S. A.* 108, 48 (Jan. 2011), 19193–19198.
- [39] David G Rand and Martin A Nowak. 2013. Human cooperation. *Trends Cogn. Sci.* 17, 8 (Aug. 2013), 413–425.
- [40] Mark Rowland, Shayegan Omidshafiei, Daniel Hennes, Will Dabney, Andrew Jaegle, Paul Muller, Julien Pérolat, and Karl Tuyls. 2021. Temporal difference and return optimism in cooperative multi-agent reinforcement learning. In *AAMAS ALA Workshop*.
- [41] Francisco C Santos, Marta D Santos, and Jorge M Pacheco. 2008. Social diversity promotes the emergence of cooperation in public goods games. *Nature* 454, 7201 (July 2008), 213–216.
- [42] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. 2016. Prioritized Experience Replay. arXiv:1511.05952 [cs.LG]
- [43] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347 (2017).
- [44] Hirokazu Shirado and Nicholas A Christakis. 2020. Network Engineering Using Autonomous Agents Increases Cooperation in Human Groups. *iScience* 23, 9 (Sept. 2020), 101438.
- [45] Hirokazu Shirado, Feng Fu, James H Fowler, and Nicholas A Christakis. 2013. Quality versus quantity of social ties in experimental cooperative networks. *Nat. Commun.* 4 (Nov. 2013), 2814.
- [46] Karl Tuyls and Ann Nowé. 2005. Evolutionary game theory and multi-agent reinforcement learning. *The Knowledge Engineering Review* 20, 1 (2005), 63–90.
- [47] Hado van Hasselt, Arthur Guez, and David Silver. 2015. Deep Reinforcement Learning with Double Q-learning. arXiv:1509.06461 [cs.LG]