

# Emergent Cooperation under Uncertain Incentive Alignment

Nicole Orzan

University of Groningen  
Groningen, The Netherlands  
n.orzan@rug.nl

Davide Grossi

University of Groningen  
Groningen, The Netherlands  
n.orzan@rug.nl

Erman Acar

ILLC & IvI, University of Amsterdam  
Amsterdam, The Netherlands  
erman.acar@uva.nl

Roxana Rădulescu

Vrije Universiteit Brussels, Belgium  
Utrecht University, The Netherlands  
roxana.radulescu@vub.be

## ABSTRACT

Understanding the emergence of cooperation in systems of computational agents is crucial for the development of effective cooperative AI. Interaction among individuals in real-world settings are often sparse and occur within a broad spectrum of incentives, which often are only partially known. In this work, we explore how cooperation can arise among reinforcement learning agents in scenarios characterised by infrequent encounters, and where agents face uncertainty about the alignment of their incentives with those of others. To do so, we train the agents under a wide spectrum of environments ranging from fully competitive, to fully cooperative, to mixed-motives. Under this type of uncertainty we study the effects of mechanisms, such as reputation and intrinsic rewards, that have been proposed in the literature to foster cooperation in mixed-motives environments. Our findings show that uncertainty substantially lowers the agents' ability to engage in cooperative behaviour, when that would be the best course of action. In this scenario, the use of effective reputation mechanisms and intrinsic rewards boosts the agents' capability to act nearly-optimally in cooperative environments, while greatly enhancing cooperation in mixed-motive environments as well.

## KEYWORDS

Multi-Agent Reinforcement Learning; Social Dilemmas; Intrinsic Rewards; Public Goods Game

### ACM Reference Format:

Nicole Orzan, Erman Acar, Davide Grossi, and Roxana Rădulescu. 2024. Emergent Cooperation under Uncertain Incentive Alignment. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 10 pages.

## 1 INTRODUCTION

The question of how cooperation can emerge in human and animal societies [20, 50], as well as in systems of artificial agents has been a focus of research for decades [5, 12, 15, 39]. Understanding the mechanisms behind cooperation emergence is fundamental for the

development of artificial agents capable of effectively collaborating with both other agents and humans. The most significant challenges in the problem of cooperation emergence are encountered in scenarios where agents' incentives exhibit some degree of misalignment, commonly known as mixed-motive environments [52]. One instance of such environments are social dilemmas, among which the public goods game [4] is the focus of the present paper.

*State of the art.* Among all the factors that can undermine cooperation emergence, uncertainty is a prominent one. There are two main types of uncertainty at play in social dilemmas: social uncertainty, which refers to uncertainty regarding the behavior of other agents [7, 18]; and environmental uncertainty, which refers to the variability of resources obtained from the environment [3, 62].

Addressing social uncertainty is the drive behind many of the mechanisms that have been studied, across disciplines, in order to foster cooperation. Direct and indirect reciprocity are among the main social mechanisms that have been demonstrated to facilitate the emergence of cooperation in scenarios with misaligned interests [53]. While direct reciprocity refers to the behaviour of acting cooperatively with individuals who previously acted cooperatively with us, indirect reciprocity refers to cooperating with individuals who behaved cooperatively with others. These cooperation-aiding mechanisms have been extensively studied in social psychology [24, 26, 55, 63], philosophy [44, 45], computational biology [43] and evolutionary game theory [40–42, 64]. The mechanism of indirect reciprocity is particularly relevant in systems where encounters among individuals are rather sparse, and agents have to rely on external information to judge others. One of the ways in which indirect reciprocity can be instantiated is by means of reputation mechanisms. Reputation is an indirect measure of the trustworthiness of an individual, usually rooted in the past conduct the individual exhibited with others [44, 45]. Research has shown how reputation can lead to faster emergence of cooperation in social dilemmas and, more generally, mixed-motive games, especially if assisted by other cooperation-driving mechanisms, such as partner selection and trust-building [1, 2, 8, 30].

To the best of our knowledge, however, it is still unclear to what extent such mechanisms remain robust even under conditions of environmental uncertainty, which has been shown to act as a detrimental factor in the provision of public goods in human groups [9, 28, 62]. In particular, we are concerned with situations in which computational agents may be uncertain about how (mis-)aligned their interests actually are in the interaction at play: *To what extent*



This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

*cooperation-aiding mechanisms remain effective under environmental uncertainty?* This is the core research question of the paper.

*Contributions.* In this work, we focus on the impact of environmental uncertainty in scenarios where agents’ incentives are (mis-) aligned to various extents. We investigate the effect of different cooperation-aiding mechanisms (reputation, steering agents and intrinsic rewards) in fostering cooperation among reinforcement learning (RL) even under this type of uncertainty.

This paper presents the three following contributions. *First*, we present the outcomes related to the impact of environmental uncertainty on independent reinforcement learning agents trained on the Extended Public Goods Game (EPGG) [46]. This game enables training across environments with a variety of incentive alignments, while introducing uncertainty on those. We find that uncertainty diminishes agents’ cooperation in cooperative and mixed-motive settings while keeping the competitive behaviour in the competitive ones. *Second*, we present an analysis of the independent role of reputation with effective social norms, steering agents and intrinsic rewards in shaping the behaviour of independent Deep Q-network (DQN) agents [37], trained on the EPGG with uncertainty. We find that the presence of an effective social norm enables agents to leverage reputation, promoting cooperative behaviour in mixed-motive and cooperative environments. This result is also achieved by the integration of intrinsic rewards based on the idea of self-play. *Third*, we analyse the entangled effect of reputation, intrinsic reward and uncertainty on the behaviour learnt by DQN agents. We show that this combination of social mechanisms brings learning agents closer to the optimal behavior in the competitive environments, particularly when steering agents are employed. Moreover, the inclusion of intrinsic rewards boosts cooperation in the cooperative environments when steering agents are not employed.

## 2 RELATED WORK

*Social Dilemmas.* Social dilemmas are group settings where the conflict between the individual and collective interest can easily lead to cooperation failure, ultimately giving rise to suboptimal outcomes for the group [17, 32]. In recent times, these scenarios are being explored using MARL [29] to gain insight into the conditions under which cooperative behaviour might emerge. In [49] a model that allows agents to optimise multiple objectives in social dilemmas with team settings is developed. In [34], with the aim of defining more complex environments than matrix games, temporally extended social dilemmas are introduced. In [31] authors work on improving communication in sequential social dilemmas setting via intrinsic motivation, by rewarding agents that have high social influence over other agents. In [22] reciprocity is employed to elicit cooperation in otherwise selfish individuals, giving rise to pro-sociality. While the aforementioned studies focus on analysing games with a mixed-motive structure, in this work we aim at investigating the impact of concurrently training over a set of environments with different levels of incentive alignments, using the public goods game as a source of inspiration.

*Reputation and Social Norms.* Reputation is a central mechanism to develop cooperation in multi-agent systems. The concept of reputation is strongly tied to the development of trust among agents

[10, 38, 47], and therefore to partner selection [2, 11, 21]. The influence of the partner selection mechanism on the emergence of cooperation among independent RL agents is investigated in [2]: agents determine their partner selection choice by observing the actions other agents took in previous interactions. Other works showed how decentralised trust can improve the consensus rate among learning agents in the presence of unreliable agents [27]. Combined effect of the mechanisms such as reputation, partner selection, and (direct and third-party) punishment on learning agents has also been studied, showing how the highest rates of cooperation are reached by these mechanisms simultaneously [16]. In [36], authors investigate intrinsic motivation for reputation in agents trained on spatially and temporally complex tasks, by directly modifying the reward signal. They observe an increase in group coordination in the game when agents are motivated to achieve a good reputation.

Studies in computational and evolutionary biology [13, 23], evolutionary game theory [35, 40–42, 58], and artificial intelligence investigated the evolution of social norms, their impact on the group dynamics, and their role in sustaining cooperation. For example, [51] examines the changes in cooperativity within a group when different social norms are applied. The findings show that the most effective norms are the ones expressing that agents should cooperate with good individuals and defect with bad ones. This is widely observed result (see [42, 48, 54], among others). In [57] author investigate the effect of public sanctioning on the emergence of social norms, to solve social dilemmas. In [1], the authors study the interaction of reputation mechanisms and social norms in a reinforcement learning framework, and show that the reputation mechanism alone is not sufficient to guide tabular RL agents towards a cooperative equilibrium. However, this approach becomes effective when combined with intrinsic rewards, or when (non-learning) agents that already adhere to the norm are introduced in the system.

In the real-world, interactions are characterised by environments that present different levels of incentive alignment. Furthermore, oftentimes it is not straightforward to determine to which extent the incentives between agents are actually aligned: agents may not be aware of whether they are playing in a cooperative, competitive, or mixed scenario. All of the aforementioned studies focus on applying social mechanisms on agents that are acting in mixed-motive environments in which the achievement of cooperative behaviour by learning is not guaranteed. However, none of these studies focuses on learning across a range of environments with different levels of incentive alignment, which is the main distinctive character of our work. We are investigating, in particular, the effect of social mechanisms – building upon the framework of [1] – in environments that exhibit different levels of incentive alignments and in presence of uncertainty. We should also note that, in contrast [1], we go beyond tabular learning approach, to accommodate the continuous nature of EPGG.

## 3 PRELIMINARIES

In this section, we introduce (1) the Extended Public Goods Game, which is the main game-theoretic setting of our work, and (2) the key social mechanisms we used to study emergence of cooperation: reputation, social norms and intrinsic rewards.

### 3.1 The Extended Public Goods Game

The Public Goods Game (PGG) is a collective decision game extensively studied in economics, which describes situations where cooperation by rational agents is Pareto optimal, but because of the profitability of free-riding, rational agents converge to defection, which is the Nash equilibrium [4]. We refer to this kind of games as mixed-motive, since the incentives of the agents are partially misaligned. We model Public Goods Games as a tuple  $\langle N, c, A, f, u \rangle$ , where  $N$  symbolises the set of players, and  $|N| = n \in \mathbb{N}$  is the number of players. Every player  $i$  is endowed with some amount of wealth (or coins)  $c_i \in \mathbb{R}^{\geq 0}$ , and  $c = (c_1, \dots, c_n)$  denotes the tuple containing all agents' coins. In this work we will be assuming all endowments to be equal, so for all  $i, j \in N$ ,  $c_i = c_j$ . Each agent can decide whether to cooperate (investing in the public good) or defect (keeping the endowment); therefore, the set  $A$  of allowed actions consists of cooperate ( $C$ ) and defect ( $D$ ) i.e.,  $A = \{C, D\}$ . The vector  $a = (a_1, \dots, a_n) \in A$  represents the action profile of the agents. The number  $f \in (1, n)$  is called the multiplication factor, and specifies the quantity by which the total good collected by the agents' investment is multiplied. This resulting quantity generates the public good, which is then evenly distributed among all agents. We define by  $u$  the vector of utilities received by the agents, and  $u_i : A^n \times (1, n) \times (\mathbb{R}^{\geq 0})^n \rightarrow \mathbb{R}$  is the utility function for agent  $i$ , specified as follows:

$$u_i(a, f, c) = \frac{1}{n} \sum_{j=1}^n c_j I(a_j) \cdot f + c_i (1 - I(a_i)), \quad (1)$$

where  $a_j$  is the  $j$ -th entry of the action profile  $a$  and  $I(a_j)$  is the indicator function, equal to 1 if the action of the agent  $j$  is cooperative, and 0 otherwise, and  $c_j$  denotes the  $j$ -th entry of  $c$ . Because we assumed  $1 < f < n$ , it is easy to check that, in PGG, the profile with all agents defecting is a dominant strategy equilibrium, while the Pareto optimal profile is the one in which all agents cooperate. This captures the social-dilemma nature of PGG. In particular, for  $n = 2$  the PGG defines a class of Prisoner Dilemma games.

Following [46], in this paper we generalise the definition of PGG by letting  $f$  range over  $(0, R_+)$  instead of  $(1, n)$ , where  $R_+ > n$  is an arbitrary value. This defines the class of Extended Public Good Games (EPGG). EPGG can model not only mixed-motives scenarios (i.e., when  $1 < f \leq n$ ) but also (1) fully competitive scenarios (i.e., when  $0 \leq f < 1$ ) in which the defection profile is a Pareto optimal dominant strategy equilibrium (thus a Nash equilibrium) and incentives are therefore fully misaligned; and (2) fully cooperative scenarios (i.e., when  $n \leq f \leq R_+$ ) in which the cooperation profile is a Pareto optimal dominant strategy equilibrium (thus an optimal Nash equilibrium) and incentives are therefore fully aligned. Note that in the boundary case  $f = 1$  both CC and DD are Pareto optimal but only DD is an equilibrium. It is important to mention already that our agents will be trained in the EPGG under uncertainty about the value of  $f$ . The uncertainty of agent  $i$  consists of an observed value of the multiplication factor  $f_i^{obs} \in (0, R_+)$  which, however, may differ from the true one (i.e.,  $f$ ). As detailed in Section 6, we model the uncertainty on the observation as Gaussian noise over the true value of  $f$  (while ensuring that any observed value  $f_i^{obs}$  less than 0 is truncated to 0).

In our work we consider interactions among two agents at a time. Figure 1 shows the normal form of the matrix games instantiating the EPGG for 2 players, with  $N = \{X, Y\}$  endowed with 4 coins each, for the following set of multiplication factors:  $f \in \{0.5, 1.0, 1.5, 3.5\}$ . These games model fully competitive ( $f \in \{0, 1\}$ ), fully cooperative ( $f \in \{3.5\}$ ) and mixed-motives ( $f = 1.5$ ) interactions.

### 3.2 Reputation

Following [1, 41], we model reputation for an agent  $i$  as a binary variable,  $r_i \in \{0, 1\}$  in which 1 (resp. 0) will denote *good*- (resp. *bad*-) *reputation*. How these values are assigned depends on *social norms*, mechanisms that drive reputation update in the group. Social norms are basically functions i.e.,  $g : A \times \{0, 1\} \rightarrow \{0, 1\}$  that update the reputation of each playing agent after every interaction. That is, the new value of the reputation  $r_i$  is computed based on the last action  $a_i$  took by the agent, and the current reputation  $r_j$  of the opponent i.e.,  $r_i = g(a_i, r_j)$ . Given two possible values for both the reputation and the available actions, there are  $2^{2^2}$  possible social norms that could be defined [41]. It has been shown that some social norms are more apt to push and maintain cooperation than others [42]. One of the most effective social norms for sustaining cooperation in mixed-motive scenarios is to assign a good reputation to those who cooperate with agents having good reputations and defect against those with bad reputations, regardless of their actions. Formally,

$$g(a_i, r_j) = \begin{cases} 1 & \text{if } a_i = C \text{ and } r_j = 1, \\ 1 & \text{if } a_i = D \text{ and } r_j = 0, \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

i.e., a reputation of 1 is assigned to an agent that cooperates with an opponent that has reputation 1, or to an agent that defects against an opponent with reputation 0, and assigns 0 otherwise.

The social norm defined in Equation (2) is the norm we deploy in our experiments. However, application of social norms to the EPGG requires one further adaptation with respect to existing approaches. Because of the uncertainty about the multiplicative factor, we need to consider the fact that agents should not be expected by the norm to cooperate when operating in fully competitive environments. To account for this, we make the social norms dependent on the multiplication factor of the game that is currently being played. So, social norms for the EPGG are functions  $G : [0, R_+] \times A \times \{0, 1\} \times \{0, 1\} \rightarrow \{0, 1\}$ . In this setting we will be working with one specific adaptation of Equation (2):

$$G(f, a_i, r_i, r_j) = \begin{cases} g(a_i, r_j) & \text{if } f \geq 1, \\ r_i & \text{otherwise} \end{cases} \quad (3)$$

That is, if the multiplication factor currently played is bigger or equal than 1 — the boundary between a competitive ( $f < 1$ ) and a mixed motive ( $f > 1$ ) interaction — we let the social norm modify the reputation of agent  $i$  as described earlier. However if the multiplication factor is less than 1, reflecting a competitive setting, the reputation of the agent stays fixed at its current value  $r_i$ . Such design is justified as, in competitive games, the cooperative action does not represent a Pareto optimal outcome nor a Nash equilibrium. Consequently, neither the group nor the individual would benefit from it. However, in the boundary case  $f = 1$  — where both full cooperation and full defection are Pareto optimal (see Figure

$f = 0.5$	Player X	$f = 1.0$	Player X	$f = 1.5$	Player X	$f = 3.5$	Player X																
	C   D		C   D		C   D		C   D																
Player Y	C <table style="display: inline-table; border-collapse: collapse;"><tr><td style="border: 1px solid black; padding: 2px 5px;">2, 2</td><td style="border: 1px solid black; padding: 2px 5px;">1, 5</td></tr><tr><td style="border: 1px solid black; padding: 2px 5px;">5, 1</td><td style="border: 1px solid black; padding: 2px 5px;">4, 4</td></tr></table>	2, 2	1, 5	5, 1	4, 4	Player Y	C <table style="display: inline-table; border-collapse: collapse;"><tr><td style="border: 1px solid black; padding: 2px 5px;">4, 4</td><td style="border: 1px solid black; padding: 2px 5px;">2, 6</td></tr><tr><td style="border: 1px solid black; padding: 2px 5px;">6, 2</td><td style="border: 1px solid black; padding: 2px 5px;">4, 4</td></tr></table>	4, 4	2, 6	6, 2	4, 4	Player Y	C <table style="display: inline-table; border-collapse: collapse;"><tr><td style="border: 1px solid black; padding: 2px 5px;">6, 6</td><td style="border: 1px solid black; padding: 2px 5px;">3, 7</td></tr><tr><td style="border: 1px solid black; padding: 2px 5px;">7, 3</td><td style="border: 1px solid black; padding: 2px 5px;">4, 4</td></tr></table>	6, 6	3, 7	7, 3	4, 4	Player Y	C <table style="display: inline-table; border-collapse: collapse;"><tr><td style="border: 1px solid black; padding: 2px 5px;">14, 14</td><td style="border: 1px solid black; padding: 2px 5px;">7, 11</td></tr><tr><td style="border: 1px solid black; padding: 2px 5px;">11, 7</td><td style="border: 1px solid black; padding: 2px 5px;">4, 4</td></tr></table>	14, 14	7, 11	11, 7	4, 4
2, 2	1, 5																						
5, 1	4, 4																						
4, 4	2, 6																						
6, 2	4, 4																						
6, 6	3, 7																						
7, 3	4, 4																						
14, 14	7, 11																						
11, 7	4, 4																						
	D <table style="display: inline-table; border-collapse: collapse;"><tr><td style="border: 1px solid black; padding: 2px 5px;">5, 1</td><td style="border: 1px solid black; padding: 2px 5px;">4, 4</td></tr></table>	5, 1	4, 4		D <table style="display: inline-table; border-collapse: collapse;"><tr><td style="border: 1px solid black; padding: 2px 5px;">6, 2</td><td style="border: 1px solid black; padding: 2px 5px;">4, 4</td></tr></table>	6, 2	4, 4		D <table style="display: inline-table; border-collapse: collapse;"><tr><td style="border: 1px solid black; padding: 2px 5px;">7, 3</td><td style="border: 1px solid black; padding: 2px 5px;">4, 4</td></tr></table>	7, 3	4, 4		D <table style="display: inline-table; border-collapse: collapse;"><tr><td style="border: 1px solid black; padding: 2px 5px;">14, 14</td><td style="border: 1px solid black; padding: 2px 5px;">7, 11</td></tr><tr><td style="border: 1px solid black; padding: 2px 5px;">11, 7</td><td style="border: 1px solid black; padding: 2px 5px;">4, 4</td></tr></table>	14, 14	7, 11	11, 7	4, 4						
5, 1	4, 4																						
6, 2	4, 4																						
7, 3	4, 4																						
14, 14	7, 11																						
11, 7	4, 4																						

**Figure 1: Normal form games instantiating the EPGG for two players (X and Y) with 4 coins each, for four possible values of the multiplication factor:  $f = 0.5$  (competitive game),  $f = 1.0$  (boundary game with both CC and DD optimal),  $f = 1.5$  (mixed-motive game) and  $f = 3.5$  (cooperative game).**

1) – the norm demands an out-of-equilibrium response attributing positive reputation to cooperation rather than defection. With the introduction of reputation, the policy of an RL agent  $i$  playing on the EPGG against an opponent  $j$  depends both on the observed value of the multiplication factor  $f_i^{obs}$ , which can be affected by uncertainty, and the reputation of the opponent  $r_j$ . Thus, the policy of agent  $i$  is a function  $\pi_i : [0, R_+] \times \{0, 1\} \rightarrow [0, 1]$ . The action taken by agent  $i$  is then sampled from the distribution:  $a_i \sim \pi_i(\cdot | f_i^{obs}, r_j)$ . Besides the social norms driving the reputation update, reputation systems may also benefit from bootstrapping effects provided by agents with fixed behaviors. In [1] authors show that relying solely on the reputation mechanism is insufficient to shift the equilibrium of the Prisoner’s Dilemma towards cooperation. However, they also noted that introducing agents that do not respond to the incentives of the game but rather exhibit fixed behaviour, strictly adhering to specific social norms, may significantly enhance cooperation among agents. They refer to such agents as *steering agents*. In our study we leverage this approach, experimenting with different proportions of steering agents in the game, to evaluate the impact of those on the learning agents. Following [1] we implement a steering agent  $i$  as an agent that follows the social norm defined by Equation (3). Noting that our objective is investigating the agent’s behaviour in the presence of observation uncertainty, we emphasise that steering agents receive uncertain observations as well. A steering agent  $i$  playing against an opponent  $j$  with reputation  $r_j$ , receiving observation  $f_i^{obs}$  on the value of the multiplication factor, employs the fixed policy  $\pi_s$ :

$$\pi_s(f_i^{obs}, r_j) = \begin{cases} 1 & \text{if } f_i^{obs} \geq 1 \text{ and } r_j = 1 \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

### 3.3 Intrinsic Reward

The idea of intrinsic reward comes from the psychological concept of intrinsic motivation [25], which characterises the drive of doing something for its inherent enjoyability rather than for any external or utilitarian gain [14]. This defines a reward signal that does not come from an external source (i.e., the environment), but from the agent itself [6]. Following the model provided by [1], we ground intrinsic rewards in the idea of self-play, namely, those are the utility that agent  $i$  imagines it would obtain when playing against itself. This choice is motivated by the goal of prioritising learning policies that are effective against agents with a similar behaviour. Intrinsic reward is defined as the utility that agent  $i$ , with reputation  $r_i$ , would receive when playing against an agent with its same reputation in an environment where the true multiplication factor is the observed one:  $u_i(\mathbf{a}', f_i^{obs}, c)$ , where  $\mathbf{a}' = (a'_i, a'_i)$  and  $a_i \sim \pi_i(\cdot | f_i^{obs}, r_i)$ .

When the intrinsic reward formulation is employed, the reward function  $R_i : A^n \times \mathbb{R}^{\geq 0} \times \mathbb{R}^n \rightarrow \mathbb{R}$  of agent  $i$  playing against agent  $j$  on a game with multiplication factor  $f$  and receiving observation  $f_i^{obs}$ , is a convex combination of the two utilities:

$$R_i = \beta u_i(\mathbf{a}, f, c) + (1 - \beta) u_i(\mathbf{a}', f_i^{obs}, c), \quad (5)$$

where  $\mathbf{a} = (a_i, a_j)$  and  $a_i \sim \pi_i(\cdot | f_i^{obs}, r_j)$ , and  $u_i(\mathbf{a}, f, c)$  is the utility received from the EPGG.  $\beta \in [0, 1]$  is a parameter that determines the weight of the game utility  $u_i(\mathbf{a}, f, c)$  against the intrinsic reward  $u_i(\mathbf{a}', f_i^{obs}, c)$  i.e., when  $\beta = 1$  (resp.  $\beta = 0$ ) the obtained reward is the game utility (resp. intrinsic reward).

## 4 METHODS

In this section we present the experimental framework we employ to investigate the effect of reputation, steering agents and intrinsic rewards on the behaviour of independent reinforcement learning agents trained on the EPGG. To assess the effect of these elements on the system, we conduct four different kind of experiments, considering both scenarios with and without uncertainty in the agents’ observation of the multiplication factor. *First*, we perform a study where the reputation and the intrinsic rewards mechanism are disabled. *Second*, we enable the reputation mechanism. In these experiments we additionally investigate the influence of the presence of steering agents to enforce norms within the system. To do so, we observe the results on the average cooperation—that is, the average relative frequency with which cooperative actions are selected by the agents—under different proportions of steering agents. *Third*, we perform experiments where only the intrinsic reward mechanism is enabled. And *fourth*, we study the interaction between reputation mechanism, steering agents and intrinsic rewards.

As described in Section 3.1, the EPGG consists of a set of games, each characterised by a multiplication factor  $f$ , that is a continuous variable in  $(0, R_+)$ . In order to compare our results to previous work on reputation mechanisms and MARL [1], we opt for tabular Q-learning [59] as a baseline. In this case, the set of multiplication factors needs to be discretised. To alleviate this restriction, we employ Deep Q-Networks (DQN) [37], as the natural function approximation extension of the tabular approach. To the best of our knowledge, ours is the first study of reputation mechanisms for emergent cooperation with DQN agents.

*Training process.* To conduct the experiments, we employ a pool of 10 agents whose endowments are all fixed to the same constant value  $c = 4$ . At every training epoch two players (that will be referred to as active agents) are uniformly sampled from the pool, and a multiplication factor  $f$  is sampled uniformly from a predefined

set of values. Afterwards, the two selected agents play repeatedly the matrix game defined by  $f$  for a fixed number of rounds  $M$ . At each of these rounds every agent  $i$  observes the current multiplication factor of the game  $f_i^{obs}$ , potentially with uncertainty (how precisely this uncertainty is modelled we will deal with in Section 6). When no uncertainty is introduced, the observed value of the  $f$  coincides with the true value:  $f_i^{obs} = f, \forall i \in N$ . If the reputation mechanism is not enabled, this observed value is the only input the active agents receive. In this case the agent’s policy takes the simplified form  $\pi_i : [0, R_+] \rightarrow [0, 1]$ , and  $a_i \sim \pi(\cdot | f_i^{obs})$ . If the reputation mechanism is enabled, the agents receive as part of the input the reputation  $r_j$  of the opponent. The type of the policy in this case is the one defined in Section 3.2.

Upon receiving their observations, the active agents take a cooperate or defect action according to their policy. After acting, the agents receive a reward from the environment, which, if the intrinsic reward mechanism is not enabled, is equal to the utility function of the EPGG for the current game being played. Otherwise, the reward function is of the form described in Section 3.3. Subsequently, the reputation of each active agent  $i$  is updated depending on the actual multiplication factor  $f$ , the action taken ( $a_i$ ), the opponent’s reputation ( $r_j$ ), in accordance with the social norm in Equation (3).

We refer to the set of EPGG games employed for training by describing the set of associated multiplication factors. In the non-tabular case (Deep Q-learning [37]), agents are trained on the continuous range of values  $F_1 = [0.5, 3.5]$ . In the tabular case (Q-learning [59]) the set  $F_2 = \{0.5, 1.0, 1.5, 3.5\}$  has been employed. This second set has been chosen because it contains games with cooperative ( $f = 3.5$ ), competitive ( $f = \{0.5, 1.0\}$ ) and mixed-motive nature ( $f = \{1.0, 1.5\}$ ), while  $F_1$  is the continuous extension of  $F_2$ .

The exploration strategy employed by the agents is  $\epsilon$ -greedy, while the evaluation policy selects greedily the action that maximises the Q-function. Agents are trained independently for 10.000 epochs, and the number of interactions between the two active agents has been fixed to  $M = 200$ . We also set the learning rate to  $\alpha = 0.01$ , discount factor  $\gamma = 0.99$  and the reputation assignment error  $\chi = 0.001$  (in accordance with [1]). When intrinsic rewards are enabled, we fixed  $\beta = 0.1$ . When using Q-learning, replicating the setup of [1], agents collect in a memory buffer all the  $M$  interactions that occur at the current epoch. The Q-learning update is implemented only after all the  $M$  interactions are completed, leaving the policy unaltered in this period. The interactions stored in the memory buffer are then discarded after the update, to account both for the possible change in the game played and the non-stationarity of the environment. For training the Q-learning algorithm, we set the current exploration rate  $\epsilon$  to 0.01. The architecture used for the DQN algorithm consists of a Multi-Layered Perceptron with one hidden layer of size 4 and a ReLU activation function. For training the DQN algorithm we employ a decaying exploration rate, with a starting value of  $\epsilon = 0.1$  and an ending value of  $\epsilon = 0.001$ . As for the Q-learning update, the DQN update takes place for the two playing agents at the end of the  $M$  interactions. The memory buffer is cleared after every interaction.

*Evaluation.* The metric we analyse is the average cooperation of the active agents after every training epoch, i.e. the average relative frequency of cooperative actions. To assess the statistical

significance between the results of the different experiments, the Welch’s t-test [60] with a p-value of 0.0001 has been employed, where the variable of interest is the average cooperation of the active agents over the last 50 training epochs of every run, also averaged over 20 runs. For both the tabular and non-tabular cases, the evaluation has been performed on the set  $F_2$ .

## 5 LEARNING WITHOUT UNCERTAINTY

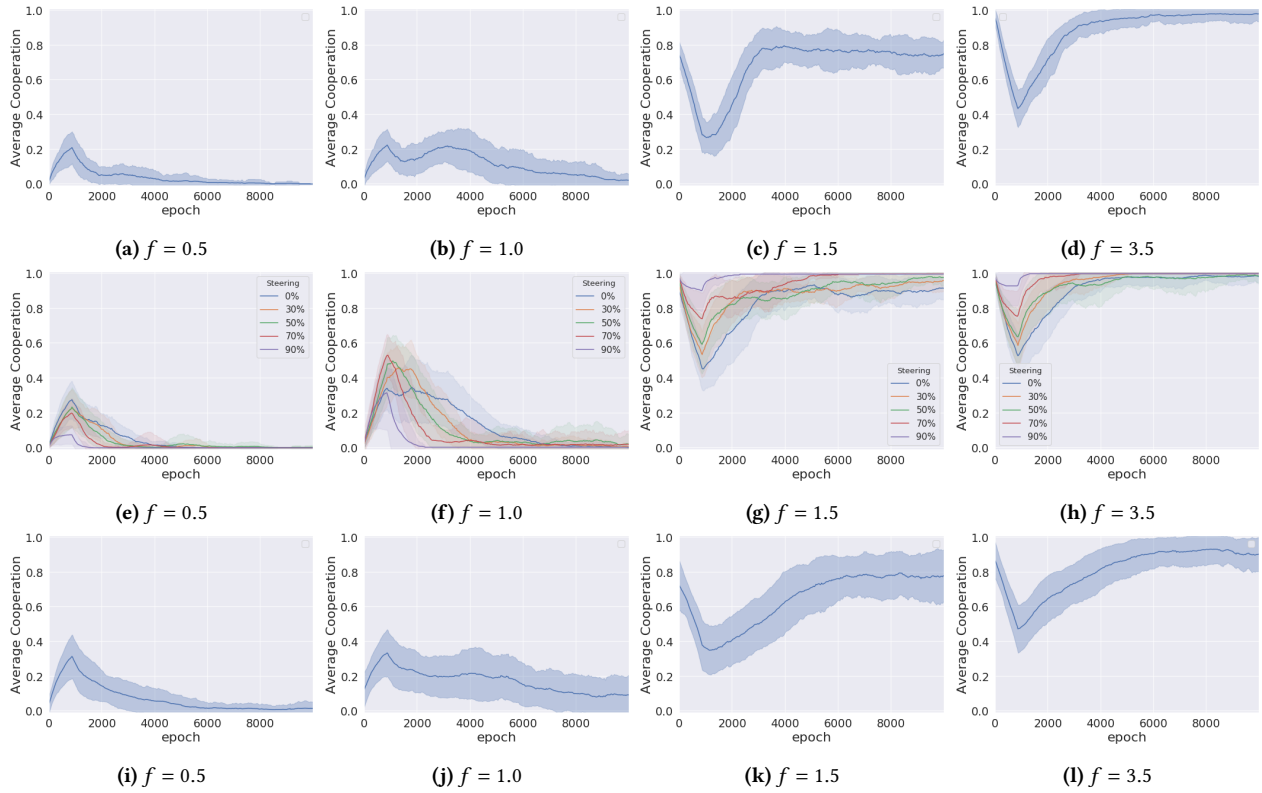
In this section, we present the baseline results of the effect of reputation, steering agents, and intrinsic rewards on the behavior of independent RL agents trained on the EPGG without uncertainty on  $f$ , so with perfect knowledge on the incentives (mis-)alignment.<sup>1</sup>

*Q-learning.* In the absence of any cooperation-aiding mechanisms, Q-learning agents quickly converge to the rational behaviour described in Section 3.1: in the cooperative game ( $f = \{3.5\}$ ) agents converge to cooperation, while in the competitive and mixed motive games ( $f = \{0.5, 1.0, 1.5\}$ ), they converge to defection. This result aligns with the baseline findings in [46], where two REINFORCE [61] agents are trained in a two-players EPGG without uncertainty. When reputation is enabled, cooperation is quickly reached and stably maintained in the cooperative environment but also in the mixed-motive one, without the need of introducing steering agents. At  $f = 0.5$  defection remains a stable equilibrium. At  $f = 1.0$  agents do not perfectly converge to defection, even if the rates of cooperation are very low. The introduction of steering agents does not significantly change the outcomes. The employment of the introspective reward mechanism has an impact only at  $m = 1.5$ : here agents do not converge to defection or cooperation, but reach an average cooperation of  $0.51 \pm 0.21$ . The combination of reputation and intrinsic rewards affects only the threshold environment  $f = 1.0$  when at least 70% of the population consists of steering agents, shifting the defective equilibria to more cooperative outcomes.

Interestingly, the result for the mixed-motive environment when employing reputation differs from the one obtained by [2] on the Prisoner’s Dilemma game: here, the authors observe that enabling reputation and effective social norms alone do not suffice to emerge and support cooperation in a system of 10 agents playing the repeated Prisoner’s Dilemma game. However, the difference of the outcome in our experiments is justifiable if we look at the gap between the utilities the agents get in the cooperation-cooperation and defection-defection equilibria: in our setting the gap among these utilities is smaller. We note that our finding is also consistent with the argument that plain tabular reinforcement learning agents can converge to cooperation in mixed-motive games even without the use of reputation and memory [19].

*Deep-Q Network.* The outcomes of the experiments for the scenario without uncertainty are depicted in Figure 2. The plots show the average cooperation of the active agents at each evaluation step, averaged over 20 runs. The results with and without reputation mechanism are consistent with those found using Q-learning in the cooperative environment  $f = 3.5$ , the competitive environment  $f = 0.5$  and the threshold one  $f = 1.0$ . However we can notice that,

<sup>1</sup>Due to space constraints, we present results and figures for the DQN algorithm, while the results for Q-learning are discussed, but the visualisation of the outcomes is available in the supplementary material.



**Figure 2: Average cooperation for DQN agents trained across environments with different multiplication factors. The top row (a - d) shows the results in the absence of cooperation-aiding mechanisms, the middle row (e - h) in the presence of reputation mechanisms and a social norm that aids cooperation, and the bottom row (i - l) in the presence of intrinsic rewards.**

even in absence of the reputation mechanism, in the mixed-motive environment  $f = 1.5$  the DQN agents reach an average cooperation value of  $0.78 \pm 0.09$ , which is significantly far from the competitive Nash equilibria of the game. The source of this outcome might be twofold. Firstly, it may arise from the fact that the DQN algorithm does not ensure convergence to the Nash equilibria. Secondly, the imbalance in interval lengths between the competitive games ( $f \in [0.5, 1.0]$ ) and the cooperative ones ( $f \in [2.0, 3.5]$ ), can be the reason leading to a greater influence of the latter on the mixed-motive environments ( $f \in [1.0, 2.0]$ ). The middle row of Figure 2 shows that the introduction of reputation, even without steering agents, statistically improves the average cooperation of the mixed-motive environment, leaving the results for the cooperative, the competitive and the threshold environments unchanged<sup>2</sup>. The incorporation of intrinsic rewards brings about the same outcome as we observe in the absence of any additional mechanisms (further details are provided in the supplementary material).

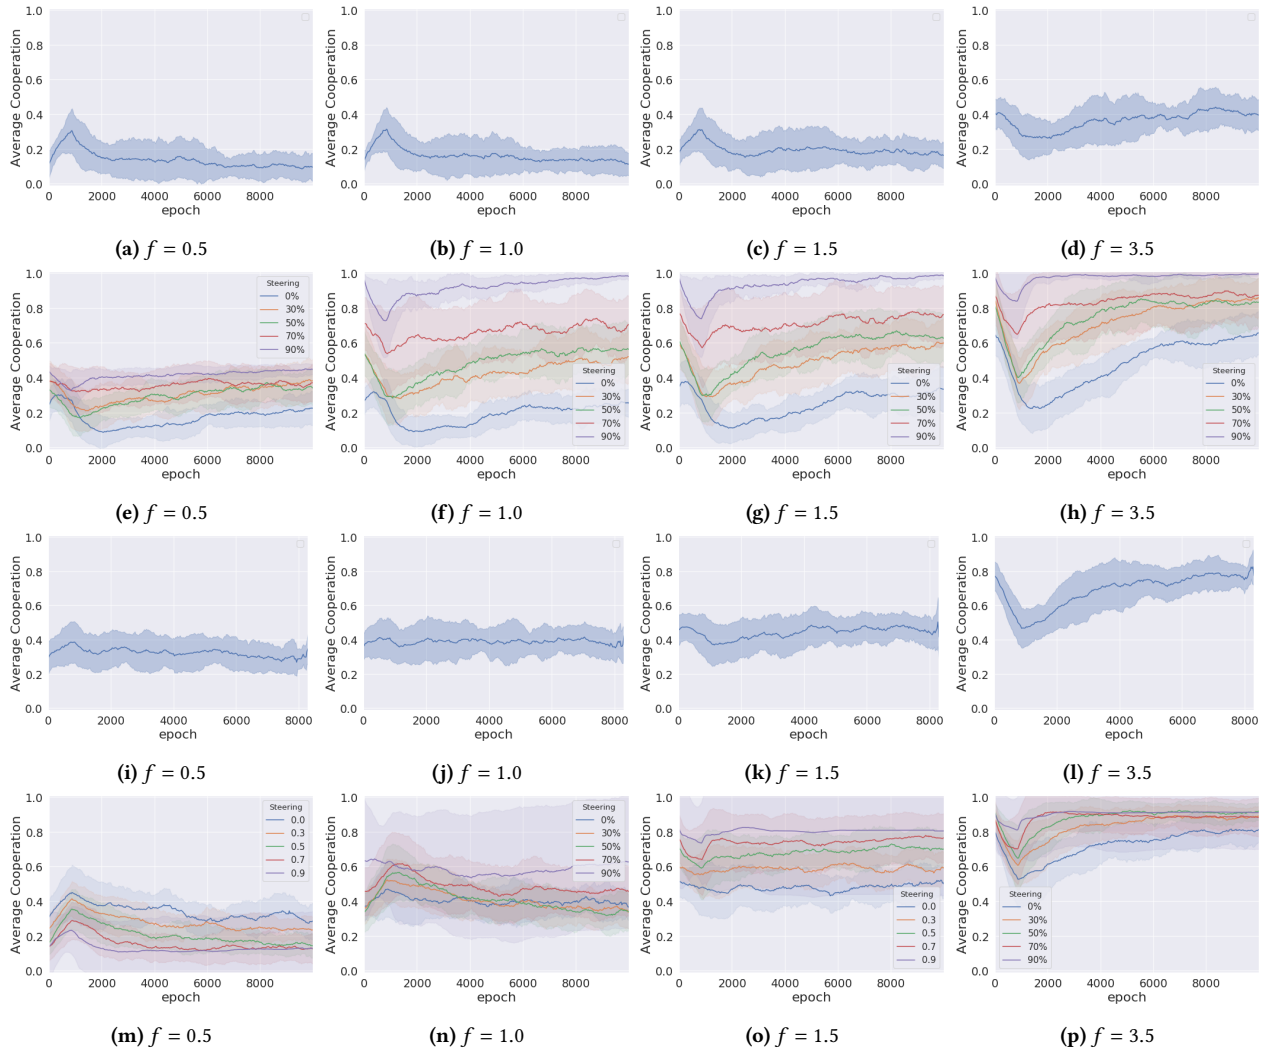
## 6 LEARNING WITH UNCERTAINTY

In this section, we investigate the role of reputation, steering agents and intrinsic rewards on the behaviour of independent RL agents trained on the EPGG when introducing uncertainty about  $f$ , so

<sup>2</sup>The statistical significance has been assessed using the Welch’s t-test [60] with a p-value of 0.0001.

with only noisy information about incentive (mis-)alignment. We model uncertainty over the observation of the multiplication factor  $f$  as Gaussian noise: for agent  $i$  the observed multiplication factor is sampled from the distribution  $f_i^{obs} \sim f + N(0, \sigma_i^2)$ , where  $\sigma_i$  is the uncertainty experienced by agent  $i$ , which is the standard deviation of the distribution. We note that, to ensure the adherence of the observed values to the definition of  $f$  inside the EPGG, we clip any sampled value  $f_i^{obs} < 0$  to 0. In these experiments, to allow RL agents to receive continuous observations we utilise the DQN algorithm. The DQN architecture and the employed hyperparameters are identical to the ones previously outlined in Section 4. Here we set all agents (both the learning agents and the steering ones) to have the same uncertainty:  $\sigma_i = \sigma, \forall i \in N$ . We fix the value  $\sigma = 2$ . In Figure 3 we present the plots for the results of the four experiments outlined in Section 4 in the presence of uncertainty. These plots depict the average cooperation for each training epoch which, we recall, is computed as an average over 20 runs. Below we provide an overview of the obtained results. For a more in-depth discussion, please refer to the supplementary material.

*Uncertainty.* The top row of Figure 3 shows the results when uncertainty is present and reputation and intrinsic rewards are not integrated into the system. When compared to results obtained without uncertainty, we note that uncertainty modifies significantly



**Figure 3: Average cooperation for DQN agents trained across environments with different multiplication factors under uncertainty ( $\sigma_i = 2 \forall i \in N$ ). Results are displayed in four rows: the first without reputation or intrinsic rewards, the second with reputation and a social norm that aids cooperation, the third with the intrinsic rewards formulation, and the fourth with reputation, a social norm that aids cooperation, and the intrinsic rewards formulation.**

**Table 1: Averages and standard deviations for the cooperation values computed for each game, in the cases without uncertainty (No U), with uncertainty (U), and with uncertainty and intrinsic rewards (UI), and p-values resulting from the Welch’s t-tests.**

Multiplication factor	No Uncertainty	Uncertainty	Uncertainty with I	p-value (No U/U)	p-value (U/UI)
0.5	0.00 ± 0.02	0.09 ± 0.07	0.31 ± 0.10	0.0021	2.5337e-8
1.0	0.02 ± 0.04	0.12 ± 0.06	0.36 ± 0.13	0.0003	1.2248e-9
1.5	0.78 ± 0.09	0.16 ± 0.06	0.45 ± 0.13	7.8726e-19	3.624e-11
3.5	0.98 ± 0.03	0.40 ± 0.07	0.78 ± 0.12	1.7306e-14	3.254e-12

the average cooperation within the cooperative and mixed-motive environments  $f \in \{1.5, 3.5\}$ , while it does not make a significant difference in the competitive and threshold ones (i.e.,  $f \in \{0.5, 1.0\}$ ). The average values and the p-values resulting from the tests are

reported in Table 1. The main effects of uncertainty as follows: 1) in the cooperative setting ( $f = 3.5$ ), agents fail to reach cooperative equilibria, causing the average cooperation to drop from  $0.98 \pm 0.03$  when uncertainty is absent, to  $0.40 \pm 0.07$  in the presence of

**Table 2: Averages and standard deviations for the cooperation in each game, averaged over the last 50 training epochs, for the experiments with uncertainty. The measures are reported separately for two scenarios: the one where only the reputation mechanism is enabled (R) and the one where both reputation mechanisms and intrinsic rewards are employed (RI). Values that exhibit statistical significance are in bold.**

$f$	Percentage of Steering Agents									
	0%		30%		50%		70%		90%	
	R	RI	R	RI	R	RI	R	RI	R	RI
0.5	0.22 ± 0.08	0.30 ± 0.11	0.29 ± 0.09	0.23 ± 0.06	<b>0.34 ± 0.09</b>	<b>0.15 ± 0.06</b>	0.37 ± 0.08	0.13 ± 0.07	0.45 ± 0.03	0.13 ± 0.20
1.0	0.25 ± 0.06	0.37 ± 0.11	0.38 ± 0.13	0.35 ± 0.08	<b>0.57 ± 0.13</b>	<b>0.33 ± 0.12</b>	<b>0.70 ± 0.15</b>	<b>0.46 ± 0.14</b>	<b>0.98 ± 0.01</b>	<b>0.62 ± 0.37</b>
1.5	0.33 ± 0.11	0.49 ± 0.12	<b>0.41 ± 0.12</b>	<b>0.59 ± 0.09</b>	0.64 ± 0.13	0.71 ± 0.14	0.76 ± 0.14	0.77 ± 0.12	0.98 ± 0.02	0.81 ± 0.36
3.5	<b>0.65 ± 0.12</b>	<b>0.83 ± 0.09</b>	<b>0.55 ± 0.09</b>	<b>0.89 ± 0.04</b>	0.83 ± 0.13	0.91 ± 0.06	0.87 ± 0.17	0.89 ± 0.10	1.00 ± 0.00	0.91 ± 0.21

uncertainty. 2) in mixed-motive games, the average cooperation drops significantly, decreasing from an average value of  $0.78 \pm 0.09$  (without uncertainty) to  $0.16 \pm 0.06$  (with uncertainty), when  $f = 1.5$ .

*Reputation.* The second row of Figure 3 shows the results of the experiments with the inclusion of the reputation mechanism and the effective social norm described in Section 3.2. We observe that, except for the competitive environment with  $f = 0.5$ , the final average cooperation presents significant differences compared to the experiments without the reputation mechanism, in particular in the absence of steering agents. We confirm the significance of the results through a t-test comparing the ‘‘Uncertainty’’ column in Table 1 with the 0% – R column in Table 2. When steering agents are introduced, the average cooperation of the system increases in all environments: just with 30% presence of steering agents, the average cooperation is significantly higher in the mixed motive and in the cooperative games (see again Tables 1 and 2). Notably, the addition of steering agents leads to a slight improvement in cooperation within the competitive scenario as well. This behavior can be explained by remembering that the observed value  $f_i^{obs}$  has a lower bound of 0 but no upper bound. Consequently, agents encounter a bigger number of observed mixed-motive and cooperative values  $f_i^{obs} \in (1, +\infty)$  rather than competitive ones  $f_i^{obs} \in (0, 1)$ .

*Intrinsic Rewards.* In this section, we describe the results of the experiments where the intrinsic reward formulation of Equation (5) is applied, which are reported in the third row of Figure 3. From the plots can be observed that, even without the addition of reputation and social norms, the average cooperation values are significantly higher in all the environments compared to the case where uncertainty is present but intrinsic rewards are not introduced (top row of Figure 3). The averages values and the p-values resulting from the tests are contained in Table 1.

*Reputation and Intrinsic Rewards.* The bottom row of Figure 3 displays the results of the combined presence of intrinsic rewards and reputation with an effective social norm. In Table 2 we compare these results with those in which the intrinsic reward is not introduced. Interestingly, for the competitive and threshold environments ( $f \in \{0.5, 1.0\}$ ), there is a significant difference among the results when at least 50% or more steering agents are included in the system, namely the combination of steering agents and intrinsic rewards lowers the average cooperation, partially recovering the defective behavior. Thus, in these cases, the formulation of

intrinsic rewards, when combined with reputation and steering agents, mitigates the optimistic bias introduced by the presence of steering agents. This effect can be observed by comparing Figures 3e and 3m. On the other hand, there are significant differences in the mixed-motive environment for 30% steering agents, and in the cooperative one for cases of 0% and 30%, namely the inclusion of intrinsic rewards in the presence of the reputation mechanism enhances cooperation in both these environments.

## 7 CONCLUSION

In this work we investigated the impact of uncertain incentive alignment on the cooperative behavior of independent reinforcement learning agents. Our aim was to assess the effectiveness of social mechanisms, such as reputation, steering agents and intrinsic rewards in promoting cooperation in environments marked by this type of uncertainty. First, we observed that uncertainty significantly lowers the agents’ cooperative behavior in the cooperative and mixed-motives games, while leaving it unchanged in the competitive and threshold ones. Adding a reputation mechanism with an effective social norm significantly boosts cooperation in the cooperative and mixed-motive environments, and without significantly impacting behavior in the competitive and threshold ones. Furthermore, steering agents induce optimism and significantly boost cooperation. The intrinsic reward formulation mimics the effect of reputation without steering agents: the two mechanisms, if introduced separately, yield similar effects. And finally, under uncertainty, the synergy between reputation and intrinsic rewards recovers the competitive behavior in the competitive environment, if a percentage of at least 50% of steering agents is employed. Simultaneously, it promotes cooperation in the mixed-motive and cooperative environments, when no steering agents are in use, or only a small fraction of them are deployed. As future work, we aim to investigate scenarios involving more than two agents, as well as more complex environments such as Cleanup and Harvest [56]. Moreover, we aim to explore the role of communication [33, 46] on social norms emergence.

## ACKNOWLEDGMENTS

This research has been supported by the Hybrid Intelligence Center, a 10-year program funded by the Dutch Ministry of Education, Culture and Science through the Netherlands Organisation for Scientific Research (NWO). Roxana Rădulescu is supported by the Research Foundation – Flanders (FWO), grant number 1286223N.



## REFERENCES

- [1] Nicolas Anastassacos, Julian García, Stephen Hailes, and Mirco Musolesi. 2021. Cooperation and Reputation Dynamics with Reinforcement Learning. In *AAMAS '21: 20th International Conference on Autonomous Agents and Multiagent Systems, Virtual Event, United Kingdom, May 3-7, 2021*, Frank Dignum, Alessio Lomuscio, Ulle Endriss, and Ann Nowé (Eds.). ACM, 115–123. <https://doi.org/10.5555/3463952.3463972>
- [2] Nicolas Anastassacos, Stephen Hailes, and Mirco Musolesi. 2020. Partner selection for the emergence of cooperation in multi-agent systems using reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 7047–7054.
- [3] Peter Andras, John Lazarus, Gilbert Roberts, and Steven J Lynden. 2006. Uncertainty and cooperation: Analytical results and a simulated agent society. *Journal of Artificial Societies and Social Simulation* (2006).
- [4] James Andreoni. 1988. Why free ride?: Strategies and learning in public goods experiments. *Journal of Public Economics* 37, 3 (1988), 291–304.
- [5] Robert Axelrod. 1981. The emergence of cooperation among egoists. *American political science review* 75, 2 (1981), 306–318.
- [6] Andrew G Barto. 2013. Intrinsic motivation and reinforcement learning. *Intrinsically motivated learning in natural and artificial systems* (2013), 17–47.
- [7] Jonathan Bendor. 1993. Uncertainty and the evolution of cooperation. *Journal of Conflict resolution* 37, 4 (1993), 709–734.
- [8] Amir Jalaly Bidgoly and Fereshteh Arabi. 2023. Robustness evaluation of trust and reputation systems using a deep reinforcement learning approach. *Computers & Operations Research* (2023), 106250.
- [9] Anders Biel and Tommy Gärling. 1995. The role of uncertainty in resource dilemmas. *Journal of Environmental Psychology* 15, 3 (1995), 221–233. [https://doi.org/10.1016/0272-4944\(95\)90005-5](https://doi.org/10.1016/0272-4944(95)90005-5) Green Psychology.
- [10] Andreas Birk. 2000. Boosting cooperation by evolving trust. *Applied Artificial Intelligence* 14, 8 (2000), 769–784.
- [11] Valerio Capraro, Francesca Giardini, Daniele Vilone, and Mario Paolucci. 2016. Partner selection supported by opaque reputation promotes cooperative behavior. *Judgment and Decision making* 11, 6 (2016), 589–600.
- [12] Damien Challet and Y-C Zhang. 1997. Emergence of cooperation and organization in an evolutionary game. *Physica A: Statistical Mechanics and its Applications* 246, 3-4 (1997), 407–418.
- [13] Fabio ACC Chalub, Francisco C Santos, and Jorge M Pacheco. 2006. The evolution of norms. *Journal of theoretical biology* 241, 2 (2006), 233–240.
- [14] Nuttapon Chentanez, Andrew Barto, and Satinder Singh. 2004. Intrinsically motivated reinforcement learning. *Advances in neural information processing systems* 17 (2004).
- [15] Allan Dafoe, Edward Hughes, Yoram Bachrach, Tantom Collins, Kevin R. McKee, Joel Z. Leibo, Kate Larson, and Thore Graepel. 2020. Open Problems in Cooperative AI. <https://doi.org/10.48550/ARXIV.2012.08630>
- [16] Nayana Dasgupta and Mirco Musolesi. 2023. Investigating the Impact of Direct Punishment on the Emergence of Cooperation in Multi-Agent Reinforcement Learning Systems. *arXiv preprint arXiv:2301.08278* (2023).
- [17] Robyn M Dawes. 1980. Social dilemmas. *Annual review of psychology* 31, 1 (1980), 169–193.
- [18] Paul Deutchman, Dorsa Amir, Matthew R Jordan, and Katherine McAuliffe. 2022. Common knowledge promotes cooperation in the threshold public goods game by reducing uncertainty. *Evolution and Human Behavior* 43, 2 (2022), 155–167.
- [19] Arthur Dolgoplov. 2022. Reinforcement Learning in a Prisoner’s Dilemma. *Available at SSRN 4240842* (2022).
- [20] Yali Dong, Boyu Zhang, and Yi Tao. 2016. The dynamics of human behavior in the public goods game with institutional incentives. *Scientific Reports* 6, 1 (2016), 28809.
- [21] Faqi Du and Feng Fu. 2011. Partner selection shapes the strategic and topological evolution of cooperation: the power of reputation transitivity. *Dynamic Games and Applications* 1 (2011), 354–369.
- [22] Tom Eccles, Edward Hughes, János Kramár, Steven Wheelwright, and Joel Z Leibo. 2019. Learning reciprocity in complex sequential social dilemmas. *arXiv preprint arXiv:1903.08082* (2019).
- [23] Paul R Ehrlich and Simon A Levin. 2005. The evolution of norms. *PLoS biology* 3, 6 (2005), e194.
- [24] Nicholas Emler. 1990. A social psychology of reputation. *European review of social psychology* 1, 1 (1990), 171–193.
- [25] Ayelet Fishbach and Kaitlin Woolley. 2022. The structure of intrinsic motivation. *Annual Review of Organizational Psychology and Organizational Behavior* 9 (2022), 339–363.
- [26] Feng Fu, Christoph Hauert, Martin A. Nowak, and Long Wang. 2008. Reputation-based partner choice promotes cooperation in social networks. *Phys. Rev. E* 78 (Aug 2008), 026117. Issue 2. <https://doi.org/10.1103/PhysRevE.78.026117>
- [27] Ho Long Fung, Victor-Alexandru Darvari, Stephen Hailes, and Mirco Musolesi. 2022. Trust-based Consensus in Multi-Agent Reinforcement Learning Systems. *arXiv preprint arXiv:2205.12880* (2022).
- [28] Lata Gangadharan and Veronika Nemes. 2009. Experimental analysis of risk and uncertainty in provisioning private and public goods. *Economic Inquiry* 47, 1 (2009), 146–164.
- [29] Sven Gronauer and Klaus Diepold. 2022. Multi-agent deep reinforcement learning: a survey. *Artificial Intelligence Review* (2022), 1–49.
- [30] Sohan Gyawali, Yi Qian, and Rose Qingyang Hu. 2021. Deep Reinforcement Learning Based Dynamic Reputation Policy in 5G Based Vehicular Communication Networks. *IEEE Transactions on Vehicular Technology* 70, 6 (2021), 6136–6146. <https://doi.org/10.1109/TVT.2021.3079379>
- [31] Natasha Jaques, Angeliki Lazaridou, Edward Hughes, Caglar Gulcehre, Pedro Ortega, DJ Strouse, Joel Z Leibo, and Nando De Freitas. 2019. Social influence as intrinsic motivation for multi-agent deep reinforcement learning. In *International conference on machine learning*. PMLR, 3040–3049.
- [32] Peter Kollock. 1998. Social Dilemmas: The Anatomy of Cooperation. *Annual Review of Sociology* 24 (1998), 183–214. <http://www.jstor.org/stable/223479>
- [33] Angeliki Lazaridou and Marco Baroni. 2020. Emergent Multi-Agent Communication in the Deep Learning Era. *CoRR abs/2006.02419* (2020). [arXiv:2006.02419](https://arxiv.org/abs/2006.02419)
- [34] Joel Z Leibo, Vinicio 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*. 464–473.
- [35] Akihiko Matsui. 1996. On cultural evolution: social norms, rational behavior, and evolutionary game theory. *Journal of the Japanese and International Economies* 10, 3 (1996), 262–294.
- [36] Kevin R. McKee, Edward Hughes, Tina O. Zhu, Martin J. Chadwick, Raphael Koster, Antonio Garcia Castaneda, Charlie Beattie, Thore Graepel, Matt Botvinick, and Joel Z. Leibo. 2023. A multi-agent reinforcement learning model of reputation and cooperation in human groups. [arXiv:2103.04982](https://arxiv.org/abs/2103.04982) [cs.MA]
- [37] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature* 518, 7540 (2015), 529–533.
- [38] Lik Mui, Mojdeh Mohtashemi, and Ari Halberstadt. 2002. A computational model of trust and reputation. In *Proceedings of the 35th annual Hawaii international conference on system sciences*. IEEE, 2431–2439.
- [39] Martin A Nowak, Akira Sasaki, Christine Taylor, and Drew Fudenberg. 2004. Emergence of cooperation and evolutionary stability in finite populations. *Nature* 428, 6983 (2004), 646–650.
- [40] Martin A Nowak and Karl Sigmund. 1998. The dynamics of indirect reciprocity. *Journal of theoretical Biology* 194, 4 (1998), 561–574.
- [41] Hisashi Ohtsuki and Yoh Iwasa. 2004. How should we define goodness?—reputation dynamics in indirect reciprocity. *Journal of theoretical biology* 231, 1 (2004), 107–120.
- [42] Hisashi Ohtsuki and Yoh Iwasa. 2006. The leading eight: social norms that can maintain cooperation by indirect reciprocity. *Journal of theoretical biology* 239, 4 (2006), 435–444.
- [43] Hisashi Ohtsuki, Yoh Iwasa, and Martin A Nowak. 2015. Reputation effects in public and private interactions. *PLoS computational biology* 11, 11 (2015), e1004527.
- [44] Gloria Origgi. 2012. A social epistemology of reputation. *Social Epistemology* 26, 3-4 (2012), 399–418.
- [45] Gloria Origgi. 2019. *Reputation: What it is and why it matters*. Princeton University Press.
- [46] Nicole Orzan, Erman Acar, Davide Grossi, and Roxana Radulescu. 2023. Emergent Cooperation and Deception in Public Good Games. <https://alaworkshop2023.github.io> 2023 Adaptive and Learning Agents Workshop at AAMAS, ALA 2023 ; Conference date: 29-05-2023 Through 30-05-2023.
- [47] Elinor Ostrom. 2003. Toward a behavioral theory linking trust, reciprocity, and reputation. *Trust and reciprocity: Interdisciplinary lessons from experimental research* 6 (2003), 19–79.
- [48] Jorge M Pacheco, Francisco C Santos, and Fabio ACC Chalub. 2006. Stern-judging: A simple, successful norm which promotes cooperation under indirect reciprocity. *PLoS computational biology* 2, 12 (2006), e178.
- [49] David Radke, Kate Larson, and Tim Brecht. 2022. The importance of credo in multiagent learning. *arXiv preprint arXiv:2204.07471* (2022).
- [50] Dustin Rubenstein and James Kealey. 2010. Cooperation, conflict, and the evolution of complex animal societies. *Nat. Educ. Knowl* 3 (2010), 78.
- [51] Fernando P Santos, Francisco C Santos, and Jorge M Pacheco. 2016. Social norms of cooperation in small-scale societies. *PLoS computational biology* 12, 1 (2016), e1004709.
- [52] Thomas C Schelling. 1958. The strategy of conflict. Prospectus for a reorientation of game theory. *Journal of Conflict Resolution* 2, 3 (1958), 203–264.
- [53] Laura Schmid, Krishnendu Chatterjee, Christian Hilbe, and Martin A Nowak. 2021. A unified framework of direct and indirect reciprocity. *Nature Human Behaviour* 5, 10 (2021), 1292–1302.
- [54] Sandip Sen and Stéphane Airiau. 2007. Emergence of Norms through Social Learning. In *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6-12, 2007*, Manuela M. Veloso (Ed.). 1507–1512. <http://ijcai.org/Proceedings/07/Papers/243.pdf>

- [55] Alex W Shaw, Vivian Li, and Kristina R Olson. 2013. Reputation is everything. (2013).
- [56] Eugene Vinitzky, Natasha Jaques, Joel Leibo, Antonio Castenada, and Edward Hughes. 2019. An Open Source Implementation of Sequential Social Dilemma Games. [https://github.com/eugenevinitzky/sequential\\_social\\_dilemma\\_games/issues/182](https://github.com/eugenevinitzky/sequential_social_dilemma_games/issues/182). GitHub repository.
- [57] Eugene Vinitzky, Raphael Köster, John P Agapiou, Edgar A Duéñez-Guzmán, Alexander S Vezhnevets, and Joel Z Leibo. 2023. A learning agent that acquires social norms from public sanctions in decentralized multi-agent settings. *Collective Intelligence* 2, 2 (2023), 26339137231162025.
- [58] Thomas Voss. 2001. *Game-theoretical perspectives on the emergence of social norms*. na.
- [59] Christopher JCH Watkins and Peter Dayan. 1992. Q-learning. *Machine learning* 8 (1992), 279–292.
- [60] Bernard L Welch. 1947. The generalization of ‘STUDENT’S’ problem when several different population variances are involved. *Biometrika* 34, 1-2 (1947), 28–35.
- [61] Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8, 3 (1992), 229–256.
- [62] Arjaan Wit and Henk Wilke. 1998. Public good provision under environmental and social uncertainty. *European journal of social psychology* 28, 2 (1998), 249–256.
- [63] Junhui Wu, Daniel Balliet, and Paul AM Van Lange. 2016. Reputation, gossip, and human cooperation. *Social and Personality Psychology Compass* 10, 6 (2016), 350–364.
- [64] Chengyi Xia, Juan Wang, Matjaž Perc, and Zhen Wang. 2023. Reputation and reciprocity. *Physics of Life Reviews* (2023).