

# Quantifying the Self-Interest Level of Markov Social Dilemmas

Richard Willis<sup>1</sup>, Yali Du<sup>1</sup>, Joel Z. Leibo<sup>1,2</sup> and Michael Luck<sup>3</sup>

<sup>1</sup>King’s College London

<sup>2</sup>Google DeepMind

<sup>3</sup>University of Sussex

richard.willis@kcl.ac.uk, yali.du@kcl.ac.uk, jzl@deepmind.com, michael.luck@sussex.ac.uk

## Abstract

This paper introduces a novel method for estimating the self-interest level of Markov social dilemmas. We extend the concept of self-interest level from normal-form games to Markov games, providing a quantitative measure of the minimum reward exchange required to align individual and collective interests. We demonstrate our method on three environments from the Melting Pot suite, representing either common-pool resources or public goods. Our results illustrate how reward exchange can enable agents to transition from selfish to collective equilibria in a Markov social dilemma. This work contributes to multi-agent reinforcement learning by providing a practical tool for analysing complex, multistep social dilemmas. Our findings offer insights into how reward structures can promote or hinder cooperation, with potential applications in areas such as mechanism design.

## 1 Introduction

Social dilemmas are situations where individual incentives conflict with group interests, presenting significant challenges in multi-agent cooperation. The main difficulty of social dilemmas is that prosocial actions can be personally costly. Agents need sufficient motivation to care about others for collective action to become more attractive than selfish behaviour. We address this with reward exchange, whereby agents agree to exchange a fixed proportion of their rewards with each other, creating an incentive for them to improve the well-being of others.

The self-interest level [Willis *et al.*, 2024] quantifies the greatest proportion of their own rewards that agents can retain while using reward exchange to resolve a social dilemma. It serves as a solution to social dilemmas and a metric for players’ propensity to cooperate, assessing the gap between individual and collective incentives. A high level indicates that players can achieve socially optimal outcomes with minimal consideration for others’ interests, while a low level indicates strong incentives for players to avoid prosocial behaviour.

While the self-interest level can be computed analytically in normal-form games, a new method is needed for more complex game structures. This paper presents a novel method

for empirically estimating the self-interest level of Markov game representations of social dilemmas using multi-agent reinforcement learning (MARL). Our approach begins by training independent agents to maximise their individual rewards, which commonly results in suboptimal collective outcomes in social dilemmas. We then introduce reward exchange and progressively increase the amount shared between agents, identifying the critical threshold where these initially selfish policies transition to collective behaviours.

Our primary contributions are twofold: we present a novel quantitative method for determining the self-interest level, superseding the qualitative approach employed in previous work [Willis and Luck, 2023]; and we provide more comprehensive experimental results on three environments featuring larger numbers of agents from the Melting Pot suite [Leibo *et al.*, 2021]. This research contributes to the growing fields of Cooperative AI [Dafoe *et al.*, 2020] and MARL in mixed-motive scenarios [Du *et al.*, 2023].

## 2 Related Work

Classically, social dilemmas have been modelled as matrix games [Axelrod, 1980; Schelling, 1973]. More recently, researchers have modelled social dilemmas using Markov games [Leibo *et al.*, 2017; Hughes *et al.*, 2018], which feature greater real-world complexities, including graded cooperativeness, agents with partial information about the state of the world, and decisions with temporally extended consequences.

Our work focuses on two prominent approaches. The first approach develops game-theoretic metrics to quantify the amount of shared interest required to achieve socially optimal equilibria in mixed-motive games [Apt and Schaefer, 2014; Elias *et al.*, 2010; Chen and Kempe, 2008; Chen *et al.*, 2011; Caragiannis *et al.*, 2010; Willis *et al.*, 2024]. These contributions are limited to analytically tractable games. Our contribution addresses this gap by extending these metrics to complex, computationally intractable games.

The second direction focuses on MARL methods for independent agents operating without centralised control or intrinsic motivations [Du *et al.*, 2023]. Within this framework, reward transfers have emerged as a mechanism for promoting collective behaviours. Baker [Baker, 2020] employs reward transfer to represent inter-agent social preferences and study their impact on team formation. Schmid *et al.* [Schmid *et*

*al.*, 2023] introduce a marketplace mechanism to trade reward shares, allowing agents to acquire stakes in each other. Gemp *et al.* [Gemp *et al.*, 2022] use reward transfers to minimise the social costs of local Nash equilibria.

Although these approaches provide valuable qualitative insights, they lack rigorous methods for determining the minimum reward transfers required to achieve socially optimal outcomes. Our work addresses this limitation by establishing a quantitative threshold for the required amount of reward exchange to resolve a social dilemma, thereby adding a descriptive dimension to the analysis and minimising transfer costs when such transfers are costly.

Several studies have explored dynamic reward transfer mechanisms via additional game actions. Lupu and Precup [Lupu and Precup, 2020] and Wang *et al.* [Wang *et al.*, 2021] investigate gift-giving actions in tragedy of the commons and coordination games, respectively. Yang *et al.* [Yang *et al.*, 2020] optimise reward transfers to shape the behaviour of learning algorithm opponents, thereby enhancing the overall reward for the transferee. Yi *et al.* [Yi *et al.*, 2022] enable agents to support collective behaviours by dynamically exchanging rewards with others nearby. With contracts, payments can be made conditional on specific joint actions being taken [Christoffersen *et al.*, 2022]. A fair value for an action has been proposed [Sodomka *et al.*, 2013].

However, action-level reward transfers require strategies to specify transfers for all joint actions across all states, increasing computational and strategic complexity. Our approach simplifies this by requiring a single advance commitment to a fixed proportion to exchange.

### 3 Background

#### 3.1 Reinforcement Learning in Markov Games

A Markov game is played by  $n$  players within a finite set of states  $\mathcal{S}$ . The game is parameterised by sets of available actions for each player  $\vec{\mathcal{A}} = (\mathcal{A}_1, \dots, \mathcal{A}_n)$ , and a stochastic transition function  $T : \mathcal{S} \times \vec{\mathcal{A}} \rightarrow \Delta(\mathcal{S})$ , mapping from joint actions at each state to the set of discrete probability distributions over states. A Markov game is partially observable if, instead of the state, the agents only view observations,  $\vec{o}_t$ , provided by the observation function  $O(s_t)$ . After a transition, each agent  $i$  receives a reward specified by their per-timestep reward function,  $r_i(\vec{a}_t, s_t)$ , where  $\vec{a}_t \in \vec{\mathcal{A}}$  is the joint action at timestep  $t$  and  $s_t \in \mathcal{S}$  is the state at timestep  $t$ . Each agent  $i$  independently learns a policy  $\pi_i(a_i|o_i)$  to maximise a long-term  $\gamma$ -discounted episode reward, defined as:

$$R_i(\vec{\pi}) = \mathbb{E}_{\vec{a}_t \sim \vec{\pi}} \left[ \sum_{t=0}^{\infty} \gamma^t r_i(\vec{a}_t, s_t) \right]$$

#### 3.2 Markov Social Dilemmas

Social dilemmas are situations in which individuals must choose between acting selfishly (to defect) for personal gain or acting in a prosocial manner (to cooperate), yielding greater overall benefits for the collective. For all agents: (i) the collective does better when an agent chooses to cooperate than when the agent chooses to defect; (ii) each

agent may be better off individually when it defects; and (iii) all agents prefer mutual cooperation over mutual defection. We assess the impact of actions on the group using a social welfare metric, which quantifies a notion of collective good. Let  $\vec{R}(\vec{\pi}) = (R_1(\vec{\pi}), \dots, R_n(\vec{\pi}))$  denote the tuple of episode rewards for all agents under joint policy  $\vec{\pi}$ . As is common in the literature [Anshelevich *et al.*, 2004; Koutsoupias and Papadimitriou, 2009; Elkind *et al.*, 2020], we use utilitarian welfare  $U$ , which measures the unweighted sum of rewards obtained by all players.

$$U(\vec{R}) = \sum_{i=1}^n R_i \quad (1)$$

An  $n$ -player Markov social dilemma is a tuple  $(M, \vec{\Pi} = \vec{\Pi}_c \cup \vec{\Pi}_d)$ , where  $M$  is a Markov game, and  $\vec{\Pi}_c$  and  $\vec{\Pi}_d$  are two disjoint sets of policies said to implement cooperation and defection, respectively. The episode rewards satisfy the following properties:

1. The utilitarian welfare is greater if an agent cooperates

$$\forall i \quad U(\vec{R}(\pi_c \frown \vec{\pi}_{-i})) > U(\vec{R}(\pi_d \frown \vec{\pi}_{-i})) \quad (2)$$

2. There is at least one joint policy profile for each agent where they are better off choosing to defect

$$\forall i \quad \exists \vec{\pi}_{-i} : R_i(\pi_d \frown \vec{\pi}_{-i}) > R_i(\pi_c \frown \vec{\pi}_{-i})$$

3. All agents prefer mutual cooperation to mutual defection

$$\forall i \quad R_i((\pi_c, \pi_c \dots \pi_c)) > R_i((\pi_d, \pi_d \dots \pi_d))$$

Where  $\vec{\pi}_{-i}$  represents the tuple of policies for all players other than player  $i$ , and  $\frown$  is a coupling operator that inserts  $\pi_i$  into  $\vec{\pi}_{-i}$  such that  $\vec{\pi} = \pi_i \frown \vec{\pi}_{-i}$ .

This technique of restricting the action space to a choice between fixed policies is known as empirical game-theoretic analysis [Walsh *et al.*, 2002; Wellman, 2006]. In practice, we use learning algorithms to discover the policies.

#### 3.3 Reward Exchange

We allow agents to enter into a contract to exchange portions of their future rewards among each other. Without reward transfers, each agent is self-interested, aiming to maximise their personal reward. Sharing rewards incentivises the agents to consider the impact of their actions on the other agents.

We introduce a parameter,  $s$ , which governs the proportion of its own rewards that an agent retains, termed the *self-interest* of the agents. The remainder,  $1 - s$ , is distributed equally among the other  $n - 1$  co-players. The post-transfer reward for agent  $i$  comprises the retained portion of its own episode reward,  $R_i$ , plus any reward received from others:

$$R'_i(\vec{R}, s) = sR_i + \frac{1-s}{n-1} \sum_{j \neq i} R_j$$

#### 3.4 Self-Interest Level

We say that a social dilemma is resolved when all agents prefer to cooperate. The self-interest level of a Markov social dilemma, denoted  $s^*$ , represents the maximum amount

of self-interest that agents can retain while using reward exchange to resolve the dilemma, and is defined as:

$$s^* = \max\{s \mid \forall i \ R'_i(\vec{R}(\pi_c \cap \bar{\pi}_{-i}), s) > R'_i(\vec{R}(\pi_d \cap \bar{\pi}_{-i}), s)\}$$

Note that when  $s = \frac{1}{n}$ , the post-transfer reward function for all agents is equivalent to maximising the utilitarian metric (Equation (1)), also referred to as a collective or team reward:

$$R'_i(\vec{R}, \frac{1}{n}) = \frac{1}{n} \sum_{j=1}^n R_j = \frac{1}{n} U(\vec{R})$$

Consequently, in a Markov social dilemma with  $s = \frac{1}{n}$ , cooperative policies are preferable to defect policies because cooperation increases the total reward due to Equation (2). However, players may still strictly prefer cooperative policies for  $s > \frac{1}{n}$ . The self-interest level, therefore, has a lower bound, and  $s^* \in [\frac{1}{n}, 1]$ .

**Example.** Consider the Public Goods Game where  $n$  players each have \$1 and can contribute any amount  $c \in [0, 1]$  to a public fund. Individual contributions up to  $c_x$  are multiplied by  $k_1$ , while those above are multiplied by  $k_2$  (where  $1 < k_2 < k_1 < n$ ). The payoff to player  $i$  contributing  $c_i$  is  $1 - c_i + \frac{1}{n} \sum_{j=1}^n (k_1 \min(c_x, c_j) + k_2 \max(0, c_j - c_x))$ . Without reward exchange, contributing nothing dominates. With reward exchange, partial cooperation ( $c = c_x$ ) emerges when  $\frac{k_2}{n} < s < \frac{k_1}{n}$ , while full cooperation ( $c = 1$ ) requires  $s < \frac{k_2}{n}$ .

## 4 Method

We present our method in two parts. First, we describe our approach to estimating the self-interest level of Markov social dilemmas. Our goal is a method that takes a Markov game and outputs its self-interest level. Then, we provide practical details, including environments and training procedures.

### 4.1 Estimating the Self-interest Level

Cooperation is a graded quantity in Markov social dilemmas, making it challenging to assess the degree to which a policy cooperates. We consider joint policies that achieve an equivalent collective reward to those trained to maximise the sum of rewards (when  $s = \frac{1}{n}$ ) to be maximally cooperating. The self-interest level is estimated to be the largest value of  $s$  for which independent policies achieve at least the same social welfare as those with a team reward, implying that they are fully cooperating.

Writing joint policies trained with a self-interest of  $s$  as  $\vec{\pi}_s$ :

$$s^* = \max_{\frac{1}{n} \leq s \leq 1} : U(\vec{R}(\vec{\pi}_s)) \approx U(\vec{R}(\vec{\pi}_{\frac{1}{n}}))$$

We wish to find strong evidence that policies trained at the self-interest level converge to cooperative equilibria, regardless of their initialisation, implying that cooperation is dominant. This is achieved by choosing challenging initialisations: policies trained without reward exchange, so the agents have incentives to act selfishly and shirk cooperation. We theorise that these policies are far from cooperative behaviours in the policy space. If, after introducing reward exchange, the policies subsequently converge to cooperative equilibria, we can be reasonably confident that cooperative equilibria are the only attractors in the policy space.

**Policy Training.** Mixed-motive environments are notoriously difficult for independent MARL [Du *et al.*, 2023], as agents treat their co-players as static, which can hinder coordination and the discovery of optimal cooperative joint policies. Furthermore, MARL can struggle due to credit assignment problems, particularly when using a team reward, which may lead to so-called ‘lazy’ agents. To mitigate these issues, we implement a curriculum learning approach [Bengio *et al.*, 2009], which helps to capture the true strategic structure of the game rather than artefacts of poor policy choices:

1. **Pretraining:** We gradually increase the number of players in the environment, starting with a single player and progressively adding players until reaching the maximum number. This allows policies to learn environmental dynamics before addressing increasingly complex multi-agent interactions.
2. **Training:** We continue training the independent policies while iteratively decreasing their self-interest. This gradual approach aims to alleviate potential catastrophic forgetting caused by the change in reward function.

**Evaluation.** Due to the stochastic nature of reinforcement learning and Markov games, we repeat the training for multiple policy initialisations. Furthermore, we identify the self-interest level with a degree of tolerance, selecting the largest value of  $s$  that achieves a total reward that is not statistically worse than the best measured:

1. Compute the mean and standard deviation of the collective reward at the end of training for each value of  $s$ , and identify  $s_{\max}$  as the  $s$  value with the largest mean.
2. Conduct a one-sided Dunnett’s test [Dunnett, 1955], a method to compare multiple samples with a single control, to assess which of the means are statistically worse than  $s_{\max}$ .
3. Choose  $s^*$  as the largest  $s$  value with a mean that is not statistically worse, otherwise  $s^* = s_{\max}$  if all are worse.

### 4.2 Environments

We evaluate our approach using three environments from the Melting Pot suite [Leibo *et al.*, 2021]: Commons Harvest, Clean Up, and Externality Mushrooms<sup>1</sup>. Commons Harvest models a tragedy of the commons scenario with a finite common-pool resource, where the challenge is to avoid overexploitation. Clean Up and Externality Mushrooms represent public goods problems, where agents can invest to improve resource quality at personal cost; the challenge here is to prevent under-provisioning. Commons Harvest and Clean Up are standard benchmarks for evaluating algorithms in social dilemmas [Jaques *et al.*, 2019; Schmid *et al.*, 2023; McKee *et al.*, 2023; Baker, 2020; Wang *et al.*, 2019], while Externality Mushrooms provides an additional test case with similar cooperative dynamics. Full details of these environments can be found in [Leibo *et al.*, 2021].

**Commons Harvest.** Commons harvest comprises seven agents that harvest apples from four large and two small apple patches. Collecting an apple provides a reward of 1. Harvested apples regrow with a probability proportional to the

<sup>1</sup>Images: Commons Harvest, Clean Up, Externality Mushrooms

number of apples remaining in the patch. However, if all the apples in a patch are harvested, it is depleted, and no apples will regrow. As long as one apple remains, given sufficient time, all the apples will regrow.

**Clean Up.** Seven agents harvest apples, each providing a reward of 1. In this environment, apples grow only if the river is sufficiently clean. Over time, the river accumulates pollution that the agents may clean. Pollution reduces the rate of apple growth, and no apples will grow if the level of pollution reaches 40% or greater.

**Externality Mushrooms.** This environment has five agents and four types of mushrooms distributed throughout the map. Red mushrooms yield a reward of 1 to the individual who eats them. Green mushrooms provide a reward of 2, shared between all agents. Blue mushrooms give a reward of 3, divided equally among all agents except the one who consumed them. Orange mushrooms destroy red mushrooms. Thus, red mushrooms represent self-interest, green mushrooms embody prosocial behaviour, blue mushrooms exemplify altruism, and orange mushrooms are a punishment mechanism. Notably, consuming mushrooms promotes the growth of new mushrooms of the same colour.

### 4.3 Implementation Details

**Learning Algorithm.** We choose Proximal Policy Optimisation (PPO) [Schulman *et al.*, 2017] for its ease of use and widespread adoption in the field. We opt against using parameter sharing (also known as self-play) to enable different agents to implement different policies. Our implementation utilises a shallow neural network containing an LSTM cell.

**Environment Configuration.** For all environments, we fix the episode length to 2000 timesteps, and we modify the observation space by compressing each grid cell from 8x8 pixels to a single pixel. In Clean Up, we scale the growth rates of apples and pollution by  $\frac{n}{7}$  when played with  $n$  players, to maintain the incentive structure. This scaling is also necessary to enable the river to be cleaned: in the default setting, pollution increases faster than a single agent can clean.

**Hyperparameter Tuning.** For each environment, we tune our hyperparameters using a single-player version, allowing us to identify the best parameters as those achieving the greatest reward, which is not necessarily the case in the multi-player case due to the mixed-motive reward structure. The exploration problem is particularly challenging in Clean Up, so we cap pollution at 40% for hyperparameter tuning. This ensures that any cleaning will cause some apples to spawn, rather than needing to clean for many consecutive timesteps. We subsequently use the best policy as the initial policy for pretraining in Clean Up, ensuring that the policies have acquired cleaning and harvesting skills, whereas the other environments start with randomly initialised policies.

**Training Specifications.** For our experiments, we use five random seeds and train for 9000 episodes (18 million environment steps) at each stage of the curriculum. We use a range of self-interest values based on the ratio of the fraction of reward an agent keeps for itself compared to the proportion of a co-player’s reward it receives, because the agents

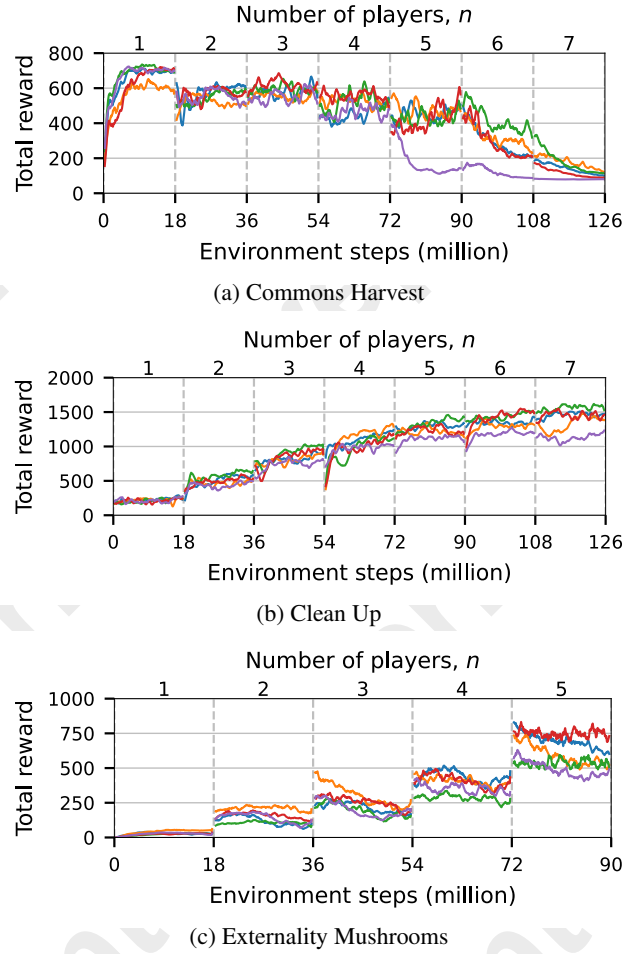


Figure 1: Pretraining, increasing numbers of players

typically face a choice between taking a benefit for themselves or allowing a co-player to gain it. The ratios we use are [20:1, 10:1, 5:1, 3:1, 5:2, 2:1, 5:3, 4:3, 1:1]. We use a p-value threshold of 0.1 for the Dunnett’s test. See [https://github.com/willis-richard/meltingpot/tree/markov\\_sd](https://github.com/willis-richard/meltingpot/tree/markov_sd) for further details.

## 5 Results

### 5.1 Pretraining

**Commons Harvest.** Figure 1a shows that the best performance is achieved when there is only one agent. In principle, multiple players should be able to match or exceed the reward of a single agent. That they do not do so in practice is due to the difficulty of coordination and the mixed-motive structure of the rewards for  $n > 1$  players. The benefits of harvesting an apple are entirely captured by the harvester, but the future cost of a reduced regrowth rate is shared among all agents; the agents may therefore have incentives to harvest more apples than is socially optimal.

Over the range of 2–4 players, all five seeds maintain good social outcomes. However, with 5–7 players, the performance collapses. At this point, the personal cost of overharvesting has decreased sufficiently to tempt agents to over-consume. This behaviour initially increases their reward, but as all poli-

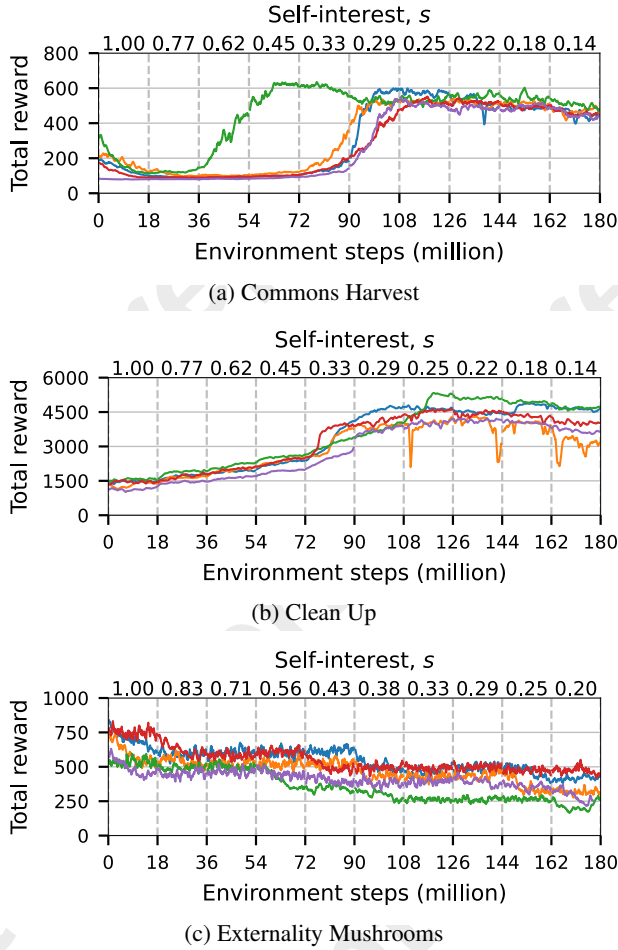


Figure 2: Iteratively decreasing self-interest during training

cies follow suit, all agents are worse off. In these equilibria, the agents quickly consume all the apples, nothing can regrow, and the tragedy of the commons materialises.

**Clean Up.** The reward increases steadily as the number of agents increases, because the apple growth rate is scaled by  $\frac{n}{7}$ . We see no collapse, as we saw for Commons Harvest, but this does not mean that the policies are achieving the maximum from the environment: they may not be optimally cleaning the river.

**Externality Mushrooms.** Figure 1c shows that social welfare increases with the number of players because harvesting mushrooms causes more to spawn; resources are not limited as in Commons Harvest. However, while social welfare increases with more agents, they are not necessarily achieving the best possible social outcomes: the agents may be under-investing in the public goods (blue mushrooms). In the next section, we examine whether reward exchange can improve the social welfare.

## 5.2 Training

Following our method from Section 4, we apply reward exchange to the pre-trained policies from Section 5.1, which serve as our challenging policy initialisations. We iteratively

s	0.46	0.33	<b>0.29</b>	0.25	0.22
mean	208	295	<b>540</b>	510	524
std dev	231	171	45	48	30
p-value	0.01	0.01	N/A	0.17	0.27

(a) Commons Harvest

s	0.33	0.29	<b>0.25</b>	0.22	0.18
mean	3646	4173	<b>4613</b>	4289	4313
std dev	416	374	439	674	504
p-value	0	0.06	N/A	0.20	0.17

(b) Clean Up

Table 1: Dunnett’s test results

decrease the self-interest of the agents to identify the critical threshold where these initially selfish policies converge to cooperative equilibria.

**Commons Harvest.** Figure 2a demonstrates that all seeds have recovered to their maximum performance by  $s = 0.29$ . Although the agents only achieve a collective reward that is less than what a single agent can achieve, this is due to the difficulty of coordination in independent MARL. In fact, by  $s = 0.14$ , the agents have an effective team reward, and they would all be better off if one of them followed the single agent policy learnt in Section 5.1 and the other six remained inactive in an empty corner.

We calculate the mean and standard deviation of the total reward achieved at the end of training for all values of  $s$ , and compute the one-sided Dunnett’s test p-value to determine whether the means are statistically lower. The results are presented in Table 1a where, for brevity, we show only the values of  $s$  close to  $s^*$ . In this case, the optimally performing value is  $s = 0.29$ . There is no greater value of  $s$  for which the mean social welfare is not statistically worse. We therefore estimate the self-interest level to be in the range  $0.29 \leq s^* < 0.33$ . Note that longer training periods or a larger number of seeds might yield slightly different results. However, limited computational resources prevent an exhaustive investigation.

**Clean Up.** Here, we also see that decreasing self-interest leads to an increase in social welfare (Figure 2b), up to a certain point. When self-interest becomes very low, one of the training seeds experiences instability issues. Table 1b displays Dunnett’s test results. The value of  $s$  with the largest mean again turns out to be the self-interest level. We therefore estimate the self-interest level to be in the range  $0.25 \leq s^* < 0.29$ .

**Externality Mushrooms.** In this environment, reward exchange does not increase social welfare, as shown in Figure 2c; performance slightly degrades as self-interest decreases. These results suggest that Externality Mushrooms does not meet our definition of a Markov social dilemma, as outlined in Section 3.2, because providing agents with incentives to care about ought to improve collective outcomes. In Section 5.4, we probe the underlying dynamics of this environment and provide further evidence that it is not a Markov social dilemma.



**Recordings.** For each environment, we provide a recording of one of the seeds at the end of pretraining and training<sup>2</sup>. For example, at the end of pretraining in Clean Up, three of the players clean the river at the beginning of the episode. When apples start spawning, they abandon their chore to harvest the apples, only returning to clean the river after too much pollution accumulates and no more apples appear. After training with reward exchange, one of these players has learnt to become a full-time cleaner. This agent harvests no apples, but still gains a proportion of the reward for each apple its co-players harvest.

### 5.3 Validation

To validate the self-interest level for Commons Harvest and Clean Up, we train new policies continuing from the  $n = 1, s = 1$  policies in Section 5.1, without using a curriculum. These policies understand the environment dynamics but have not encountered other agents before. We compare policies trained with  $s = 1$  (fully independent),  $s = s^*$  (self-interest level),  $s = \frac{1}{n}$  (team reward) and a value of  $s$  slightly larger than the range within which  $s^*$  was determined, which we call  $s^+$ .

Our results in Figure 3 confirm that the policies do not converge to cooperative equilibria without reward exchange. The lines show the mean reward received across the seeds, and the shaded regions represent a 95% confidence interval. Training at the self-interest level or with a team reward reaches a cooperative equilibrium, as expected. For both environments, the best performance is achieved with  $s^+$ . We posit that training at the game’s self-interest level offers an advantage: while socially optimal policies are learnt, the agents retain a greater amount of self-interest than when using a team reward. This larger individual incentive may alleviate credit assignment difficulties and potential lazy agent issues associated with team rewards.

We expect it to be easier to converge to cooperative equilibria when starting training from single-agent policies, because we do not begin in any particular equilibrium, compared to starting from a selfish equilibrium. As demonstrated by the  $s^+$  runs, which also reach cooperative equilibria, albeit more slowly, it is possible to train cooperative agents with a self-interest level greater than  $s^*$ . However, the possibility of converging to a selfish equilibrium may remain. Therefore, we recommend using the self-interest level to ensure policies can escape such outcomes, as demonstrated in Section 5.2.

### 5.4 Assessment

To assess whether our environments are Markov social dilemmas, we examine whether they satisfy the inequalities outlined in Section 3.2. We select sets of cooperative and defection policies and evaluate the performance of different policy combinations. Due to the symmetry of our environments, as the agents have homogeneous capabilities and randomised starting locations, the salient point is the average reward received by defecting and cooperating agents, given the number of agents choosing cooperative policies.

<sup>2</sup>[https://github.com/willis-richard/meltingpot/tree/markov\\_sd/videos](https://github.com/willis-richard/meltingpot/tree/markov_sd/videos)




Figure 3: Training without curriculum learning

We take  $n_c$  cooperators from  $\Pi_c$  and  $n - n_c$  defectors from  $\Pi_d$  and evaluate the mean rewards that defect policies and cooperate policies achieve over 225 episodes, for  $n_c = 0, 1, \dots, n$ . We plot these results as a Schelling diagram [Schelling, 1973], which shows the mean reward to an additional agent depending on whether it chooses to cooperate or to defect, given the number of co-players cooperating. On the right axis, we also plot the social welfare as a function of the number of cooperators.

**Commons Harvest.** We select one of our validation (Section 5.3) seeds with  $s = 1$  to represent the defect policies, and one with  $s = s^*$  to represent the cooperate policies. The former typically harvest the last apple in a patch, while the latter are more restrained. The Schelling diagram is displayed in Figure 4a.

We conclude that Commons Harvest is a Markov social dilemma because: (1) the social welfare strictly increases as the number of cooperative policies increases; (2) an agent always benefits from choosing a defecting policy compared to a cooperating policy, regardless of how many co-players have chosen to defect; and (3) all agents prefer mutual cooperation to mutual defection, which follows from the game symmetry.

**Clean Up.** Again, we select one of our validation seeds with  $s = 1$  to represent defect policies, and one with  $s = s^*$  to represent cooperate policies. Both the defect policies and cooperate policies tend to clean the river, but the cooperate policies typically start cleaning at a lower pollution level.

Clean Up Figure 4b is also a Markov social dilemma. It takes at least two agents simultaneously cleaning the river to reduce the pollution level, which is why we see a jump in social welfare as we move from 1 to 2 cooperators.

**Externality Mushrooms.** We select one of our pretraining seeds (Section 5.1) to represent cooperative policies, as these achieve the highest social welfare. For defect policies, we

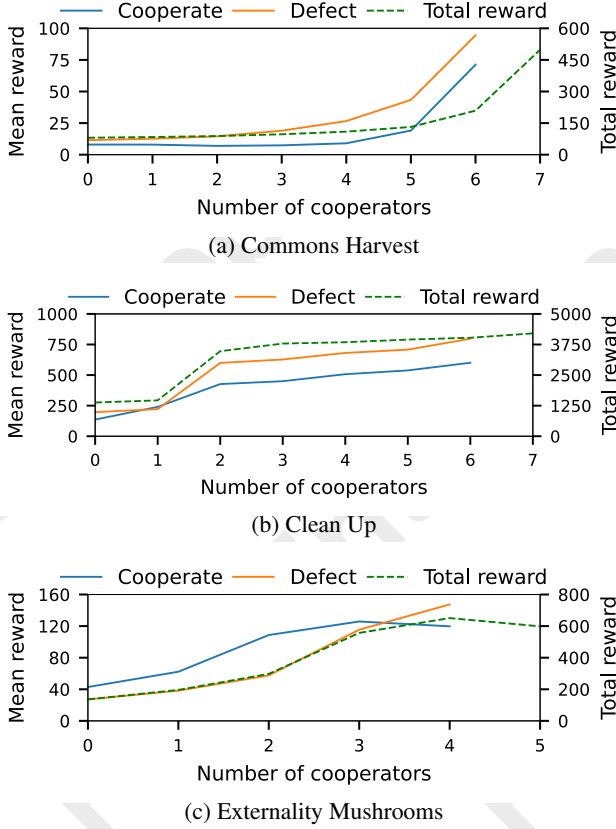


Figure 4: Schelling diagrams

train new policies with  $s = 1$  in the environment, without curriculum learning. The former policies have a greater tendency to harvest prosocial mushrooms compared to the latter.

The Schelling diagram in Figure 4c reveals that this is not a Markov social dilemma because social welfare is maximised when one player defects, violating Equation (2). This helps explain why reward exchange does not improve performance in Section 5.2: exchanging rewards with co-players likely makes it more challenging for an agent to learn the optimal selfish strategy. In general, the agents are better off choosing cooperative policies. This result was not immediately apparent from the environment design, as it appeared to incentivise agents to prioritise individual rewards over collective benefits.

Our analysis reveals that harvesting green (prosocial) mushrooms is strategically optimal even for purely self-interested agents. Although green mushrooms initially provide lower individual rewards than red (selfish) mushrooms, because harvesting a particular mushroom causes more of the same variety to grow, agents ultimately benefit from their co-players’ subsequent increased consumption of prosocial mushrooms.

PPO fails to discover this strategic advantage without curriculum learning because the temporal credit assignment challenge is significant: rewards from co-players’ subsequent prosocial mushroom consumption occur many timesteps later. Consequently, the policies almost exclusively harvest

selfish mushrooms, converging to an equilibrium with low individual rewards, although agents would individually benefit from switching to cooperative policies. The discrepancy between our cooperative and defect policies arises because prosocial mushrooms reward agents with  $\frac{2}{n}$ , which is larger when  $n$  is smaller. Only with curriculum learning, progressively adding more players, does PPO learn to appreciate the long-term benefits of harvesting green mushrooms.

## 6 Conclusion

We introduced a novel method for estimating the self-interest level of Markov social dilemmas, bridging the gap between game-theoretic metrics and complex multi-agent reinforcement learning models. The self-interest level is a valuable metric for assessing the propensity to cooperate in mixed-motive games by quantifying the gap between individual and collective incentives.

Applying our method to environments from the Melting Pot suite, we determined which environments represent genuine Markov social dilemmas, estimating their self-interest level, and which do not. We showed that policies trained with reward exchange at the self-interest level converge to equilibria with good social welfare in Markov social dilemmas (Section 5.3).

Our work offers both practical metrics and solutions for real-world social dilemmas. As a metric, the self-interest level enables risk assessment: systems with low self-interest levels face significant barriers to cooperation and may be prone to conflict, allowing system designers to identify where intervention is necessary. Additionally, assessing the self-interest level can provide insights into why certain reinforcement learning algorithms perform better or struggle in specific environments, aiding practitioners in selecting appropriate algorithms for mixed-motive scenarios. As a solution mechanism, reward exchange can be applied to problems like fishery management, where traditional quotas often fail due to persistent incentives to overfish. If fishing nations were to exchange a proportion of their fishing profits with other fishing nations, this would reduce each country’s incentive to overexploit fish stocks while simultaneously motivating all participating countries to improve ocean health.

Future work in this area could include developing a method to determine the general self-interest level [Willis *et al.*, 2024] for Markov social dilemmas. This generalisation allows agents greater freedom in their reward transfers, which can improve efficiency in asymmetrical environments. For example, if there are multiple types of agents, we could use separate parameters to govern the proportion of reward exchanged between agents of different types. Furthermore, the experimental approach in this paper could be applied to more environments, such as those in the Melting Pot suite, to delve into their underlying dynamics and challenges. We could investigate the impact on the self-interest level of systematic changes to an environment, for example, by increasing resource scarcity or introducing a sanctioning mechanism.

## Acknowledgments

This work was supported by UK Research and Innovation [grant number EP/S023356/1], in the UKRI Centre for Doctoral Training in Safe and Trusted Artificial Intelligence (www.safeandtrustedai.org) and a BT/EPSRC funded iCASE Studentship [grant number EP/T517380/1].

Compute resources were provided by King’s College London [King’s College London e-Research team, 2024].

## References

- [Anshelevich *et al.*, 2004] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The Price of Stability for Network Design with Fair Cost Allocation. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 295–304, Rome, Italy, 2004. IEEE.
- [Apt and Schaefer, 2014] Krzysztof R. Apt and Guido Schaefer. Selfishness Level of Strategic Games. *Journal of Artificial Intelligence Research*, 49:207–240, February 2014.
- [Axelrod, 1980] Robert Axelrod. Effective Choice in the Prisoner’s Dilemma. *Journal of Conflict Resolution*, 24(1):3–25, March 1980.
- [Baker, 2020] Bowen Baker. Emergent Reciprocity and Team Formation from Randomized Uncertain Social Preferences. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020*, virtual, December 2020.
- [Bengio *et al.*, 2009] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 41–48, Montreal Quebec Canada, June 2009. ACM.
- [Caragiannis *et al.*, 2010] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, Maria Kyropoulou, and Evi Papaioannou. The Impact of Altruism on the Efficiency of Atomic Congestion Games. In Martin Wirsing, Martin Hofmann, and Axel Rauschmayer, editors, *Trustworthy Global Computing*, volume 6084, pages 172–188. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- [Chen and Kempe, 2008] Po-An Chen and David Kempe. Altruism, selfishness, and spite in traffic routing. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 140–149, Chicago IL USA, July 2008. ACM.
- [Chen *et al.*, 2011] Po-An Chen, Bart De Keijzer, David Kempe, and Guido Schäfer. The Robust Price of Anarchy of Altruistic Games. In Ning Chen, Edith Elkind, and Elias Koutsoupias, editors, *Internet and Network Economics*, volume 7090, pages 383–390. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.
- [Christoffersen *et al.*, 2022] Phillip J. K. Christoffersen, Andreas A. Haupt, and Dylan Hadfield-Menell. Get It in Writing: Formal Contracts Mitigate Social Dilemmas in Multi-Agent RL, 2022.
- [Dafoe *et al.*, 2020] Allan Dafoe, Edward Hughes, Yoram Bachrach, Tatum Collins, Kevin R. McKee, Joel Z. Leibo, Kate Larson, and Thore Graepel. Open Problems in Cooperative AI, December 2020.
- [Du *et al.*, 2023] Yali Du, Joel Z. Leibo, Usman Islam, Richard Willis, and Peter Sunehag. A Review of Cooperation in Multi-agent Learning, December 2023.
- [Dunnett, 1955] Charles W Dunnett. A multiple comparison procedure for comparing several treatments with a control. *Journal of the American Statistical Association*, 50(272):1096–1121, 1955.
- [Elias *et al.*, 2010] Jocelyne Elias, Fabio Martignon, Konstantin Avrachenkov, and Giovanni Neglia. Socially-Aware Network Design Games. In *2010 Proceedings IEEE INFOCOM*, pages 1–5, San Diego, CA, USA, March 2010. IEEE.
- [Elkind *et al.*, 2020] Edith Elkind, Angelo Fanelli, and Michele Flammini. Price of Pareto Optimality in hedonic games. *Artificial Intelligence*, 288:103357, November 2020.
- [Gemp *et al.*, 2022] Ian Gemp, Kevin R. McKee, Richard Everett, Edgar A. Dueñez-Guzmán, Yoram Bachrach, David Balduzzi, and Andrea Tacchetti. D3C: Reducing the Price of Anarchy in Multi-Agent Learning. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, pages 498–506, Online, February 2022. International Foundation for Autonomous Agents and Multiagent Systems.
- [Hughes *et al.*, 2018] Edward Hughes, Joel Z. Leibo, Matthew Phillips, Karl Tuyls, Edgar Dueñez-Guzmán, Antonio García Castañeda, Iain Dunning, Tina Zhu, Kevin McKee, Raphael Koster, Heather Roff, and Thore Graepel. Inequity aversion improves cooperation in intertemporal social dilemmas. In *32nd Conference on Neural Information Processing Systems*, pages 3330–3340, Montréal, Canada, 2018. Curran Associates, Inc.
- [Jaques *et al.*, 2019] Natasha Jaques, Angeliki Lazaridou, Edward Hughes, Caglar Gulcehre, Pedro A. Ortega, D. J. Strouse, Joel Z. Leibo, and Nando de Freitas. Social Influence as Intrinsic Motivation for Multi-Agent Deep Reinforcement Learning. In *Proceedings of the 36th International Conference on Machine Learning*, pages 3040–3049. PMLR, June 2019.
- [King’s College London e-Research team, 2024] King’s College London e-Research team. King’s Computational Research, Engineering and Technology Environment (CREATE), 2024.
- [Koutsoupias and Papadimitriou, 2009] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, May 2009.
- [Leibo *et al.*, 2017] Joel Z. Leibo, Vinicius Zambaldi, and Marc Lanctot. Multi-agent Reinforcement Learning in Sequential Social Dilemmas. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems*, pages 464–473, São Paulo, Brazil, May 2017. ACM.



- [Leibo *et al.*, 2021] Joel Z. Leibo, Edgar Duéñez-Guzmán, Alexander Sasha Vezhnevets, John P. Agapiou, Peter Sunehag, Raphael Koster, Jayd Matyas, Charles Beattie, Igor Mordatch, and Thore Graepel. Scalable Evaluation of Multi-Agent Reinforcement Learning with Melting Pot. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, pages 6187–6199. PMLR, July 2021.
- [Lupu and Precup, 2020] Andrei Lupu and Doina Precup. Gifting in Multi-Agent Reinforcement Learning. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems*, pages 789–797, Auckland, New Zealand, May 2020. International Foundation for Autonomous Agents and Multiagent Systems.
- [McKee *et al.*, 2023] 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. A multi-agent reinforcement learning model of reputation and cooperation in human groups, February 2023.
- [Schelling, 1973] Thomas C. Schelling. Hockey Helmets, Concealed Weapons, and Daylight Saving: A Study of Binary Choices With Externalities. *Journal of Conflict Resolution*, 17(3):381–428, September 1973.
- [Schmid *et al.*, 2023] Kyrill Schmid, Michael Kölle, and Tim Matheis. Learning to Participate through Trading of Reward Shares. In *Proceedings of the 15th International Conference on Agents and Artificial Intelligence*, volume 1, pages 355–362, Lisbon, Portugal, February 2023. SCITEPRESS.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal Policy Optimization Algorithms, August 2017.
- [Sodomka *et al.*, 2013] Eric Sodomka, Elizabeth M Hilliard, Michael L Littman, and Amy Greenwald. Coco-Q: Learning in Stochastic Games with Side Payments. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of 3, pages 1471–1479, Atlanta, Georgia, USA, June 2013. JMLR.org.
- [Walsh *et al.*, 2002] William E Walsh, Rajarshi Das, Gerald Tesauro, and Jeffrey O Kephart. Analyzing Complex Strategic Interactions in Multi-Agent Systems. In *AAAI-02 Workshop on Game-Theoretic and Decision-Theoretic Agents*, pages 109–118, June 2002.
- [Wang *et al.*, 2019] Weixun Wang, Jianye Hao, Yixi Wang, and Matthew Taylor. Achieving cooperation through deep multiagent reinforcement learning in sequential prisoner’s dilemmas. In *Proceedings of the First International Conference on Distributed Artificial Intelligence*, pages 11:1–11:7, Beijing China, October 2019. ACM.
- [Wang *et al.*, 2021] Woodrow Z. Wang, Mark Beliaev, Erdem Biyık, Daniel A. Lazar, Ramtin Pedarsani, and Dorsa Sadigh. Emergent Prosociality in Multi-Agent Games Through Gifting. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pages 434–442, Montreal, Canada, 2021. ijcai.org.
- [Wellman, 2006] Michael P Wellman. Methods for Empirical Game-Theoretic Analysis. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference*, pages 1552–1556, Boston, Massachusetts, USA, July 2006. AAAI Press.
- [Willis and Luck, 2023] Richard Willis and Michael Luck. Resolving social dilemmas through reward transfer commitments. In *Proceedings of the Adaptive and Learning Agents Workshop*, London, May 2023.
- [Willis *et al.*, 2024] Richard Willis, Yali Du, Joel Z. Leibo, and Michael Luck. Resolving social dilemmas with minimal reward transfer. *Autonomous Agents and Multi-Agent Systems*, 38(2):49, October 2024.
- [Yang *et al.*, 2020] Jiachen Yang, Ang Li, Mehrdad Farajtabar, Peter Sunehag, Edward Hughes, and Hongyuan Zha. Learning to Incentivize Other Learning Agents. In *Proceedings of the 34th Conference on Neural Information Processing Systems*, volume 33, pages 15208–15219, Vancouver, Canada, 2020. Curran Associates, Inc.
- [Yi *et al.*, 2022] Yuxuan Yi, Ge Li, Yaowei Wang, and Zongqing Lu. Learning to Share in Multi-Agent Reinforcement Learning. In *Proceedings of the 36th Conference on Neural Information Processing Systems*. Curran Associates Inc., 2022.