

Fairness over Equality: Correcting Social Incentives in Asymmetric Sequential Social Dilemmas

Alper Demir
Izmir University of Economics
& University of Edinburgh
Izmir, Turkey
alper.demir@ieu.edu.tr

Hüseyin Aydın
Middle East Technical University
& Utrecht University
Ankara, Turkey
huseyin@ceng.metu.edu.tr

Kale-ab Abebe Tessera
University of Edinburgh
Edinburgh, United Kingdom
k.tessera@ed.ac.uk

David Abel
University of Edinburgh
Edinburgh, United Kingdom
david.abel@ed.ac.uk

Stefano V. Albrecht
DeepFlow
London, United Kingdom
stefano.albrecht@deepflow.com

ABSTRACT

Sequential Social Dilemmas (SSDs) provide a key framework for studying how cooperation emerges when individual incentives conflict with collective welfare. In Multi-Agent Reinforcement Learning, these problems are often addressed by incorporating intrinsic drives that encourage prosocial or fair behavior. However, most existing methods assume that agents face identical incentives in the dilemma and require continuous access to global information about other agents to assess fairness. In this work, we introduce asymmetric variants of well-known SSD environments and examine how natural differences between agents influence cooperation dynamics. Our findings reveal that existing fairness-based methods struggle to adapt under asymmetric conditions by enforcing raw equality that wrongfully incentivize defection. To address this, we propose three modifications: (i) redefining fairness by accounting for agents' reward ranges, (ii) introducing an agent-based weighting mechanism to better handle inherent asymmetries, and (iii) localizing social feedback to make the methods effective under partial observability without requiring global information sharing. Experimental results show that in asymmetric scenarios, our method fosters faster emergence of cooperative policies compared to existing approaches, without sacrificing scalability or practicality.

KEYWORDS

Asymmetry, Sequential Social Dilemmas, Multi-Agent Reinforcement Learning

ACM Reference Format:

Alper Demir, Hüseyin Aydın, Kale-ab Abebe Tessera, David Abel, and Stefano V. Albrecht. 2026. Fairness over Equality: Correcting Social Incentives in Asymmetric Sequential Social Dilemmas. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 12 pages. <https://doi.org/10.65109/CZPZ7833>



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

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/CZPZ7833>

1 INTRODUCTION

Multi-Agent Reinforcement Learning (MARL) addresses problems involving multiple agents, each with potentially distinct preferences and policies [2]. The dynamics of these interactions require agents to navigate not only the environment but also the actions of their counterparts, rendering MARL problems highly non-stationary [21]. This complexity is notably amplified when navigating situations where individual and group objectives diverge. Social dilemmas such as Prisoner's Dilemma are examples of such games where agents' individual benefit are in tension with the groups' overall reward. In MARL, these types of problems are temporally extended, leading to Sequential Social Dilemmas (SSDs) [16]. The literature has mainly proposed solutions centered around cultivating intrinsic motivations that drive agents toward prosocial behavior, thereby promoting actions conducive to group welfare [9, 11, 24]. Alternatively, studies explore agents incentivizing one another toward prosocial behavior [10, 19, 34].

Fairness-based intrinsic motivation methods [11, 20] have shown promise in encouraging cooperation in multi-agent systems. These approaches typically assume that agents are identical, sharing the same capabilities, roles, and reward structures, and thus interpret differences in rewards as differences in cooperative behavior. While this assumption holds in symmetric environments such as the classic Prisoner's Dilemma, it breaks down in more realistic, asymmetric settings in which agents are not identical. Under such asymmetries, even identifying Pareto-optimal joint outcomes becomes non-trivial, since the feasible payoff surface itself differs across agents with unequal reward ranges [7].

In *asymmetric* SSDs, agents may differ in physical capabilities, access to environmental resources, or reward structures, leading to unequal opportunities and outcomes. This disparity can come about even when agents behave identically, highlighting the need for fairness methods that handle such asymmetries. For example, consider the design of a tax system. Requiring every individual to pay the same flat amount may appear equal but could be seen as inherently unfair, as it places a disproportionately heavy burden on those with lower incomes. A more fair approach accounts for individuals' capacity to pay—such as through progressive taxation, where those who earn more contribute more. Beyond income, individuals with greater influence on the broader society are often expected to contribute more to the common good, whether through higher taxes,

leadership responsibilities, or adherence to social norms. Inspired by these principles, we argue that fairness in multi-agent systems should also account for asymmetries among agents.

Furthermore, while prior work [8, 11, 20, 25, 32] considers environments with partial observability, they often assume that agents can see the chosen actions and received rewards of other agents. In particular, [11] and [20] assume that agents can observe the rewards of others in order to make comparisons, while [25] instead uses the aggregate reward of all agents and allocates a portion of it back to each agent. This limits their applicability in realistic scenarios, where such information about others is often inaccessible (e.g. during inference).

In this study, we make four key contributions. **First**, we introduce novel asymmetric versions of well-known SSDs. These variants make it possible to analyze how natural differences between agents influence the dynamics of cooperation and how fairness-driven methods respond to such asymmetry. **Second**, we revise the definition of fairness in SSDs by accounting for agents’ potential reward ranges, enabling more meaningful comparisons across agents with different capabilities and better cooperation. **Third**, we propose an agent-based weighting mechanism for intrinsic social term, ensuring that agents with greater impact on the environment are held to higher standards of prosociality, thereby better addressing inherent asymmetries. **Finally**, we modify existing methods to operate effectively under partial observability by incorporating local social feedback, eliminating the need for direct access to other agents’ internal information.

Our results show existing fairness-based methods struggle with SSD asymmetries. In contrast, our approach ensures equitable social comparisons, preventing misleading incentives, while agent-based weighting improves environmental adaptation. Finally, local social feedback matches global-access performance without requiring centralized information.

2 BACKGROUND & RELATED WORK

Social Dilemmas play an important part in the field of multi-agent systems as they present the conflicts of defective and cooperative behaviors reflecting the trade-off between individual benefits and collective interests. Prisoner’s Dilemma [31] is a well-known example of a social dilemma that presents these core characteristics. Several studies model moral decision-making across various dilemmas and analyze the impact of social norms on the matrix form of the dilemmas [29, 30].

While most studies primarily assume that agents are identical in terms of their potential actions and the rewards associated with cooperative or defection strategies, there are few studies deviating from these assumptions and exploring scenarios like the asymmetric Prisoner’s Dilemma [3, 4, 6, 28]. In asymmetric setups, various agent types emerge, where a subset of agents may experience a stronger dilemma compared to others. Under such a stronger dilemma, the advantage from defecting against a cooperator or the prospect of being exploited can be larger for a subset of agents. An advantaged agent may get more out of mutual cooperation, or a disadvantaged agent might need to pay an additional cost for cooperation [35]. Such asymmetries exist in real life [33]. Further studies examine the impact of asymmetry in social dilemmas, especially

in Prisoner’s Dilemma [3, 4, 6, 28]. They show that asymmetry causes a decrease in the rate of cooperation and makes cooperation difficult to sustain unless it is actively addressed. This can be due to the propensity of disadvantaged agents to defect more frequently compared to their advantaged counterparts [3, 4, 12, 28]. It has been reported that fairness is a requirement for cooperative behavior but its definition varies between individuals, especially in asymmetric cases [12]. Lopez et al. [18] and Zhang et al. [35] both underscore how asymmetry reduces cooperation, with lower cooperative behavior observed in groups with unequal rewards. Lopez et al. note that high-reward agents are particularly less likely to cooperate, while Zhang et al. find a significantly higher cooperation rate in symmetric reward groups. Leibo et al. [16] present *Sequential Social Dilemmas* (SSDs), a temporal extension of a social dilemma. For examining such problems, four metrics such as efficiency, equality, sustainability and peace are introduced [23] which respectively measure: (i) the total reward per step, (ii) how rewards are distributed among agents, (iii) the extent to which positive rewards are sustained throughout an episode, and (iv) the degree to which agent policies exhibit defecting behavior. With rational agents and without social incentives, a phenomenon called the *tragedy of the commons* occurs where low sustainability is observed. Köster et al. [14] further identify start-up and free-rider problems in a setting where agents have individualistic preferences.

Solutions to avoid the tragedy and improve group performance in SSDs consider prosocial agents that take selfless actions. There are studies that propose a parameter to control an agent’s prosociality [24], an adjustment to selfishness level to induce cooperation [25], intrinsic drive based on other agent’s models to calculate social impact [13] and evolving intrinsic motivation for social preferences [9]. Alternatively, several studies allow agents to give rewards to others with the aim of incentivizing cooperation [19, 34]. Finally, recent approaches have explored mechanisms that go beyond intrinsic motivation. Christoffersen et al. [8] introduces formal contracting, where agents exchange rewards based on whether their actions are deemed cooperative or defective under a predefined contract. Another approach [32] encourages agents to learn to disapprove of actions in contexts where disapproval is likely to come from others, promoting socially aligned behavior through learned social feedback.

These proposed solutions often depend on global access to agents’ rewards or actions, which poses a significant limitation. For fairness, agents are typically motivated to act prosocially by comparing their rewards or actions with those of others [11, 20, 25], requiring centralized information. Similarly, classifying behaviors as cooperative or defective often requires a centralized critic [8, 32]. These approaches conflict with the assumption of independent learning under partial observability. Furthermore, potential asymmetry in these social dilemmas is rarely addressed in the literature, where agents have different action spaces or reward functions.

Table 1: The matrix form of a social dilemma game between the row agent i and the column agent j .

	C	D
C	R_i, R_j	S_i, T_j
D	T_i, S_j	P_i, P_j

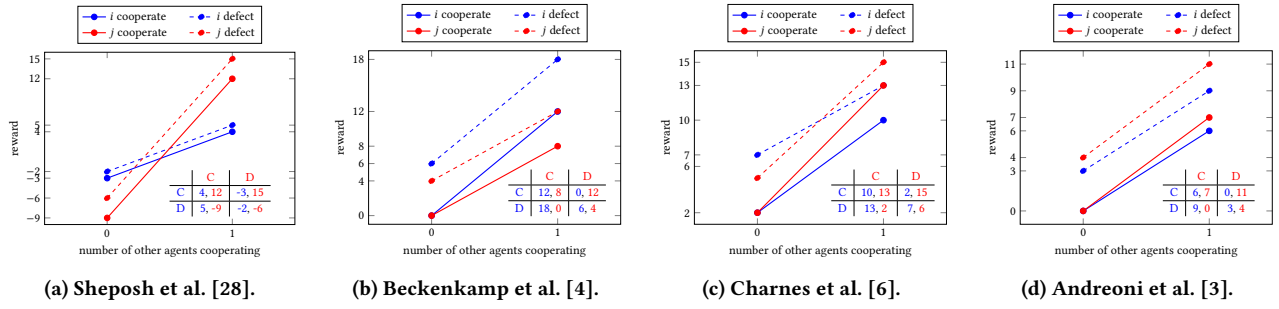


Figure 1: Schelling diagrams for example asymmetric Prisoner’s Dilemma games from the literature, between agent i (in blue) and agent j (in red). For each agent, defective strategies lead to higher rewards.

3 PRELIMINARIES

Table 1 presents a generic matrix game between agents i and j representing a social dilemma, where R_i and R_j denote the rewards when both choose to cooperate, T_i and T_j denote the rewards when they defect against a cooperating opponent, S_i and S_j denote when they cooperate while their opponent defects, P_i and P_j denote the rewards when both choose to defect.

A social dilemma in this form must satisfy all of the following conditions [16, 25]:

- C1:** $R_i > P_i$ and $R_j > P_j$: Each agent prefers mutual cooperation to mutual defection,
- C2:** $R_i > S_i$ and $R_j > S_j$: Each agent prefers mutual cooperation to being exploited by a defector,
- C3:** $2R_i > T_i + S_i$ and $2R_j > T_j + S_j$: Each agent prefers mutual cooperation to alternating between cooperation and defection,
- C4:** Either
 - C4.a:** *Greed*: $T_i > R_i$ and $T_j > R_j$: there is a temptation to defect against a cooperator,
 - C4.b:** or *Fear*: $P_i > S_i$ and $P_j > S_j$: defection in face of being exploited is preferred.

Prisoner’s Dilemma like social dilemmas involve both *greed* and *fear* factors satisfying C4.a and C4.b at the same time. If only the *greed* factor is present, it is known as a *Chicken Game*, and if only the *fear* factor is present, it is called a *Stag Hunt Game* [2, 25].

DEFINITION 1 (SEQUENTIAL SOCIAL DILEMMA). A *sequential social dilemma (SSD)* is a tuple $(\mathcal{M}, \Pi^C, \Pi^D)$ where \mathcal{M} is a partially observable stochastic game such that there exist states $s \in S$ for which the induced empirical reward matrix satisfies the social dilemma inequalities (C1-C4), and Π^C and Π^D are disjoint sets of policies that are said to implement cooperation and defection, respectively.

Such a game \mathcal{M} with N -agents is defined [2] as a tuple $\langle \mathcal{N}, S, \{A_i\}_{i \in \mathcal{N}}, T, \{R_i\}_{i \in \mathcal{N}}, \Omega, \{O_i\}_{i \in \mathcal{N}} \rangle$ where $\mathcal{N} = \{1, 2, \dots, N\}$ is the set of agents, S is the finite set of states, A is the joint set of actions such that $A = A_1 \times A_2 \times \dots \times A_N$ (A_i is the finite set of actions for agent i), $\mathcal{T} : S \times A \rightarrow \Delta(S)$ is the transition function with $\Delta(S)$ as the set of discrete probability distributions over S , $\mathcal{R}_i : S \times A \times S \rightarrow \mathbb{R}$ defines the reward function for agent i , $\Omega : S \times A \rightarrow \Delta(O)$ specifies the observation function, $O = O_1 \times O_2 \times \dots \times O_N$ is the joint set of observations so that the agent i receives a local observation $o_i \in O_i$.

Agents learn a joint policy $\pi = (\pi_1, \dots, \pi_N)$ where each agent i aims to learn its own policy π_i that maximizes its expected return given the policies of the others π_{-i} , $U_i(\pi_i, \pi_{-i}) = \mathbb{E}_{s^0 \sim \mu} [V_i^\pi(s^0)]$ where $V_i^\pi(s) = \mathbb{E}_{\pi, \mathcal{T}} [\sum_{t=0}^{\infty} \gamma^t R_i(s^t, a^t, s^{t+1}) \mid s^0 = s]$ and μ is the initial state distribution with $\gamma \in [0, 1]$ as the discount factor, $s^t \in S$ the state and $a^t \in A$ the joint action at time step t .

3.1 Base Methods

Inequity Aversion (IA) [11] and Social Value Orientation (SVO) [20] are used in this study as representative fairness-based intrinsic motivation methods to evaluate how our approach better addresses the fairness problem in asymmetric settings. Both methods provide an intrinsic drive for fairness that takes into account the temporal feature of the problem environments, calculate the temporally smoothed reward of an agent i at time step t as

$$e_i^t = \gamma \cdot \lambda \cdot e_i^{t-1} + r_i^t, \tag{1}$$

where r_i^t is the reward taken at time step t , $e_i^0 = 0$, γ is the discount factor and λ is a hyperparameter controlling the decay rate. Temporally smoothed reward of an agent is an indication of how well the agent is doing within an episode and it is a base point for comparison.

Using these smoothed rewards, IA modifies the reward of the agent i at time step t as

$$\begin{aligned} \hat{r}_i^t &= r_i^t \\ &- \frac{\alpha}{N-1} \cdot \sum_{j \neq i} \max(e_j^t - e_i^t, 0) \\ &- \frac{\beta}{N-1} \cdot \sum_{j \neq i} \max(e_i^t - e_j^t, 0), \end{aligned} \tag{2}$$

where α and β are parameters to control *disadvantageous* and *advantageous* inequity, respectively. A higher α penalizes the agent when it receives less reward than others, while a higher β penalizes when it receives more.

Similarly, using the temporally smoothed rewards, SVO calculates a reward angle for agent i at time step t as

$$\theta_i^t = \text{atan}\left(\frac{e_{-i}^t}{e_i^t}\right), \tag{3}$$

where $e_{-i}^t = \frac{1}{N-1} \sum_{j \neq i} e_j^t$. The method requires a target social value orientation θ_i^{SVO} for an agent i and modifies the reward of the agent

as

$$\hat{r}_i^t = r_i^t - w \cdot |\theta_i^{SVO} - \theta_i^t|, \quad (4)$$

with w as a weight to control the effect of SVO. Having $\theta^{SVO} = 45^\circ$ creates an incentive to be equal with the rest of the group in terms of rewards.

3.2 Asymmetry in Social Dilemmas

Social dilemmas can feature various asymmetries between agents. These asymmetries may arise from differences in available actions (abilities) or rewards (payoffs). Our analysis focuses on asymmetries that can be reflected in rewards and potentially alter agent behaviors.

We define asymmetry in a social dilemma as a situation in which agents differ in how they experience the outcomes of key strategic interactions. Specifically, a social dilemma is *asymmetric* if there is any inequality in the reward values assigned to each agent: mutual cooperation yields different rewards ($R_i \neq R_j$); the temptation to defect against a cooperator varies between agents ($T_i \neq T_j$); the cost of being exploited while cooperating differs ($S_i \neq S_j$); or mutual defection results in unequal rewards ($P_i \neq P_j$). Although all agents face the same social dilemma, differences in rewards under identical situations create unequal incentives, which can lead to different experiences and behaviors within the same environment.

Fig. 1 presents several examples of asymmetric Prisoner’s Dilemma games from the literature. In all cases, the defining inequalities of the social dilemma are preserved. That is, in each case, defection yields a higher reward than cooperation for any individual agent. However, the introduced asymmetry leads to different characteristics across the examples. While asymmetry is present in all cases of the games, the advantage shifts depending on the context.

In Fig. 1b and Fig. 1d, one agent consistently receives equal or greater rewards than the other, regardless of the outcome. In contrast, Fig. 1a shows a more nuanced dynamic: agent i (blue) earns more than agent j (red) when facing a defecting opponent, but less when the opponent cooperates. A similar role reversal appears in Fig. 1c, though it only occurs under mutual defection—agent i benefits more than agent j in that case, whereas the advantage flips in mutual cooperation.

These examples illustrate that asymmetry in reward structure does not always imply a consistent advantage for a particular agent across all strategic scenarios.

Extending from matrix games to more complex problems, based on Leibo et al.’s definition [16] of an SSD, in this general setting, we can define the rewards of an agent i for a state s satisfying social dilemma conditions as:

$$\begin{aligned} R_i(s) &= V_i^\pi(s) \text{ when } \pi_i \in \Pi^C, \pi_{-i} \in \Pi^C, \\ T_i(s) &= V_i^\pi(s) \text{ when } \pi_i \in \Pi^D, \pi_{-i} \in \Pi^C, \\ S_i(s) &= V_i^\pi(s) \text{ when } \pi_i \in \Pi^C, \pi_{-i} \in \Pi^D, \\ P_i(s) &= V_i^\pi(s) \text{ when } \pi_i \in \Pi^D, \pi_{-i} \in \Pi^D. \end{aligned}$$

DEFINITION 2 (ASYMMETRIC SEQUENTIAL SOCIAL DILEMMA). *An SSD becomes asymmetric if and only if there exist agents $i, j \in \mathcal{N}$ with $i \neq j$ and a state $s \in \mathcal{S}$ such that at least one of $R_i(s) \neq R_j(s)$, $T_i(s) \neq T_j(s)$, $S_i(s) \neq S_j(s)$, or $P_i(s) \neq P_j(s)$ holds and the terms satisfy the conditions C1 to C4.*

Table 2: Example asymmetric Prisoner’s Dilemma game between the row agent i (blue) and the column agent j (red) before (a) and after normalization (b).

	(a) original		(b) after normalization	
	C	D	C	D
C	-1, -6	-3, -5	0.67, 0.67	0, 1
D	0, -8	-2, -7	1, 0	0.33, 0.33

4 PROPOSED APPROACH

Existing methods like IA and SVO promote cooperation by comparing agents’ rewards (often coupling intrinsic and extrinsic signals in a single update [26]), implicitly assuming that differences reflect variations in cooperative behavior—as in symmetric games such as the Prisoner’s Dilemma. In asymmetric settings, however, identical policies can yield different rewards due to inherent differences in roles or capabilities. Equality-driven methods may misinterpret these natural differences as unfairness, penalizing agents who are fully cooperating (Eqs. 2 and 4). As a result, such methods can even incentivize harmful behavior, as lower-reward agents may defect to reduce the gap, undermining both cooperation and fairness. The key issue is that raw reward comparisons ignore underlying asymmetries.

4.1 Adjusting for Agent Potential

To fairly evaluate cooperation or defection, we propose that agents must be assessed relative to their own reward scales. That is, an agent’s position in its own outcome range, between its lowest and highest possible rewards ($S_i < P_i < R_i < T_i$), should be considered. Since this range can differ across agents, comparing their raw rewards directly is misleading. Instead, each agent’s behavior should be located within its own cooperative-defective range before any cross-agent comparison is made.

Consider the asymmetric Prisoner’s Dilemma shown in Table 2, where the rewards represent the number of years each prisoner is sentenced. In this scenario, the agent j (red) has a prior conviction, resulting in an additional 5 years added to its sentence for any outcome. Despite the asymmetry, the social dilemma conditions C1 to C4 are still satisfied for both agents: that is, $S_i < P_i < R_i < T_i$ and $S_j < P_j < R_j < T_j$. However, in this example, the temptation reward T_j of agent j (red) when defecting against a cooperating agent is numerically smaller than the sucker reward S_i of agent i (blue) in the same case. This is allowed under the dilemma inequalities but can lead to misleading interpretations. A direct comparison between T_j and S_i might suggest that agent j is more cooperative than agent i even though agent j is defecting and agent i is cooperating. To address this issue, a normalization step is applied, resulting in the transformed game shown in Table 2. In the normalized version, each agent’s behavior is expressed relative to their own reward range, enabling a more accurate comparison of the relative advantages and disadvantages between agents.

As a general approach, we propose a normalization step to adjust agents’ rewards during comparison over time in an SSD. Both IA and SVO utilize temporally smoothed rewards as a measure of how well one agent is doing at a time step. A *normalized* temporally

smoothed reward is defined as

$$\hat{e}_i^t = \frac{e_i^t - e_i^{\min}}{e_i^{\max} - e_i^{\min}}, \quad (5)$$

where e_i^t is the temporally smoothed reward as defined in Eq. 1 and $e_i^{\max} = \max_t(e_i^0, \dots, e_i^t)$ and $e_i^{\min} = \min_t(e_i^0, \dots, e_i^t)$. This normalization does not affect the agents' extrinsic rewards, but is used exclusively within the IA and SVO social terms. With this modification, IA and SVO become fairer and better suited to asymmetric settings, using the normalized temporally smoothed rewards in Eqs. 2 and 3.

This modification provides several improvements. First, it reduces the unintended incentive to punish cooperative agents who happen to earn higher rewards, helping to maintain the focus on comparing cooperative versus defective policies. As a result, agents are evaluated based on how well they perform relative to their own potential. Second, the normalization step ensures that values remain bounded and non-negative, providing greater control for IA and preventing negative reward angles in SVO. In environments with negative rewards, SVO can produce negative angles, which may be ambiguous. A negative value can arise in two distinct scenarios: (1) agent i receives a positive reward while others receive negative rewards, or (2) agent i receives a negative reward while others receive positive rewards. In the first case, agent i is acting selfishly; in the second, it is being highly altruistic. Despite reflecting opposite behaviors, both cases may yield the same negative angle in SVO. Normalization eliminates such ambiguity by preventing negative values, leading to more interpretable and consistent measurements.

4.2 Agent-Based Weighting of Social Influence

To better address asymmetry, differences in agents' potential impact on the environment should be reflected in their commitment to social preferences. Agents with inherent advantages must be held to a stricter standard of social responsibility. This can be implemented using agent-specific weighting by introducing a term, the *social drive modifier* as we call it, into the Eqs. 2 and 4.

Hence, we extend the definition for the weights of disadvantageous and advantageous inequity for agent i in IA as $\bar{\alpha}_i = \phi_i \cdot \alpha$ and $\bar{\beta}_i = \phi_i \cdot \beta$, where $\phi_i \geq 0$ is the social drive modifier of agent i . Similarly, in SVO, the weight of agent i is redefined as $\bar{w}_i = \phi_i \cdot w$. With these adjustments, an agent with a larger potential of influencing the environment (same action causing a bigger effect) can have a higher intrinsic drive to be prosocial. While the original parameters such as α , β and w in IA and SVO determine the importance of the social term in the reward functions in general, ϕ allows agent-specific adjustment to address inherent differences between agents, selected based on the environment and the degree of asymmetry between agents.

4.3 Localizing Social Feedback

Finally, we propose a decentralized approach in order to make IA and SVO consistent with the partially observable nature of these settings as defined in Def. 1. In this approach, each agent maintains local estimates of other agents' temporally smoothed rewards, which are updated through direct communication when available.

Algorithm 1 Update on Temporally Smoothed Rewards

```

1: for each agent  $i$  do
2:    $\triangleright$  Update estimates for non-visible agents
3:   for each agent  $k \notin \mathcal{N}_i^t$  do
4:      $j^* \leftarrow \arg \max_{j \in \mathcal{N}_i^t} \tau_{j,k}^{t-1}$ 
5:      $\hat{e}_{i,k}^t \leftarrow \hat{e}_{j^*,k}^{t-1}$ 
6:      $\tau_{i,k}^t \leftarrow \tau_{j^*,k}^{t-1}$ 
7:   end for
8:    $\triangleright$  Update own temporally smoothed reward
9:    $e_i^t \leftarrow \gamma \cdot \lambda \cdot e_i^{t-1} + r_i^t$ 
10:   $e_i^{\max} \leftarrow \max_t(e_i^0, \dots, e_i^t)$ 
11:   $e_i^{\min} \leftarrow \min_t(e_i^0, \dots, e_i^t)$ 
12:   $\hat{e}_i^t \leftarrow \frac{e_i^t - e_i^{\min}}{e_i^{\max} - e_i^{\min}}$ 
13:   $\triangleright$  Update estimates for visible agents
14:  for each agent  $j \in \mathcal{N}_i^t$  do
15:     $e_{i,j}^t \leftarrow e_j^t$ 
16:     $\tau_{i,j}^t \leftarrow t$ 
17:  end for
18: end for

```

Let $\hat{e}_{i,j}^t$ be agent i 's estimate at time t of agent j 's normalized temporally smoothed reward. Let $\tau_{i,j}^t$ denote the time step at which $\hat{e}_{i,j}^t$ was last updated, and let \mathcal{N}_i^t be the set of other agents visible to agent i at time step t . After each time step t , the temporally smoothed rewards and agents' local estimates are updated according to Alg. 1. After these steps, each agent i uses its own estimates for other agents in the Eqs. 2 and 3 where e_j^t is replaced by $e_{i,j}^t$.

This localized mechanism relaxes the strong assumption made by methods like IA and SVO, which typically require full access to all agents' temporally smoothed rewards at every time step. By adjusting to partial observability, these methods are adapted to more realistic settings where agents have limited local information. In this setup, an agent can directly observe how well a visible agent is performing and, through communication, indirectly obtain information about how non-visible agents were performing in the past via the estimates held by other visible agents.

These three modifications make the Fair&Local versions of IA and SVO better suited to asymmetric SSDs under partial observability. As raw equality can incentivize defection, the objective is fairness through mutual cooperation, allowing agents to achieve the best outcomes within their respective ranges. With this approach, agents learn cooperative policies more effectively without requiring global information.

5 EXPERIMENTS

Our experiments study the following questions: (1) How does asymmetry in sequential social dilemmas affect learning, both with and without equality-based intrinsic motivation? (2) Does our proposed method, by accounting for agent potential and applying agent-based weighting, mitigate the challenges introduced by asymmetry in terms of cooperation? (3) Can our local-feedback variants match the performance of their global-access counterparts?

Table 3: Agent types used in the Coins and Harvest. In Coins, “mismatch coin” refers to collecting the other agent’s coin.

Agent Type	Reward per Item	Abilities / Effects	Environment(s)
Standard	+1.0	Movement; zap in Harvest only	Harvest, Coins
Low-reward	+0.5	Movement; in Coins, collecting a mismatch coin penalizes the other by -3.0	Harvest, Coins
High-reward	+1.5	Movement; in Coins, collecting a mismatch coin penalizes the other by -1.0	Harvest, Coins
Wide-zap	+1.0	Larger zapping radius	Harvest
Spawn-biased	+1.0	Collecting a mismatch coin makes only this agent’s coins spawn for 25 steps	Coins




Figure 2: (a) Coins environment with 2 standard (1 red and 1 blue) agents; (b) Harvest environment with 10 standard agents.

5.1 Environments

We experiment using two well-known benchmark environments regarding Sequential Social Dilemmas ¹.

The Coins environment [9, 17], shown in Fig. 2a, involves two agents, red and blue. Coins of both colors spawn at a given rate. Collecting any coin yields +1 reward, but picking up the other agent’s coin penalizes that agent by -2 . A cooperative agent must therefore avoid collecting mismatch coins.

The Harvest environment [23, 24], shown in Fig. 2b, features 10 agents collecting apples, each worth +1. Apple regrowth depends on leaving nearby apples uncollected; over-harvesting depletes resources. Agents can also zap others with a beam, temporarily removing them from the environment for 25 steps. The dilemma arises because quickly collecting apples maximizes individual reward but undermines long-term group benefit.

Unlike matrix-form social dilemmas, SSDs do not have fixed cooperate or defect actions; instead, these labels must be assigned to behaviors. In Harvest, cooperation means avoiding apple collection in low-density areas, while in Coins it means collecting only one’s own coins. In Harvest, zapping non-defective agents or harvesting the last remaining apples are considered defective behaviors.

To study the role of asymmetry, we use variations of the environments where agents are divided into two subgroups with different characteristics. Their social dilemma properties are verified using Schelling diagrams in Appendix A.

5.1.1 Asymmetry in rewards. Here, base rewards are scaled: low-reward agents receive half the original reward, while high-reward agents receive one and a half times. In Coins, penalties are also scaled: when a low-reward agent collects the coin of a high-reward agent, the latter incurs a penalty of -3 , while the reverse case results in only -1 . These multipliers were chosen to make asymmetry visible without destabilizing training. These changes yield two variants: Coins with one high- and one low-reward agent (*asymmetry*

in coin rewards), and Harvest with five of each (*asymmetry in apple rewards*).

5.1.2 Asymmetry in actions. Here, the mechanics of actions differ between subgroups. In Harvest, wide-zap agents have an extended laser radius. In Coins, the spawn-biased agent causes only its own coins to appear for 25 steps after collecting the other’s coin. Rewards remain unchanged; only the effects of actions differ. These yield two additional variants: Coins with one standard and one spawn-biased agent (*asymmetry in coin spawn*), and Harvest with five standard and five wide-zap agents (*asymmetry in zap radius*).

Table 3 summarizes the agent types used in our experiments.

5.2 Baselines

Agents are trained with Independent DQN (IQL) [22]. We use three baselines: plain IQL to show the direct effect of asymmetry, IQL with IA (Eq. 2) to test inequity aversion’s limits, and IQL with SVO (Eq. 4) to assess social value orientation’s ability to handle asymmetry. Our proposed methods, Fair&LocalIA and Fair&LocalSVO, extend these approaches with three modifications, enabling fairer adaptation to asymmetric environments. The implementation details and used hyperparameters are given in Appendix B.

5.3 Evaluation Metrics

The experiments are evaluated using several metrics that capture different aspects of agent behavior.

Let $G_i = \sum_{t=0}^T r_i^t$ denote the return of agent i in an episode of T time steps, where r_i^t is the reward received by agent i at time step t . The *average return* is defined as $\hat{G} = \frac{1}{N} \sum_{i=1}^N G_i$, where higher values indicate better overall learning performance.

In Coins, the prosocial behavior can be measured by *the proportion of own coins* as the ratio of number of own collected coins to the number of total collected coins, averaged over agents.

Having the zapping action and a different dynamic, Harvest requires other metrics. *Sustainability* measures how evenly positive rewards are distributed over time and is defined as $S = \frac{1}{N} \sum_{i=1}^N t_i$, where $t_i = \mathbb{E}[t \mid r_i^t > 0]$ represents the expected time step at which agent i receives positive reward. Lower values suggest that most rewards were collected early in the episode. *Peace* captures the level of conflict among agents and is defined as $P = N - \frac{1}{T} \sum_{i=1}^N \sum_{t=1}^T \mathbb{I}\{\text{agent } i \text{ timed out on step } t\}$, where $P \in [0, N]$, and lower values indicate more aggression in the environment. *The number of other agents zapped* by agent i , is denoted as Z_i , with average $\bar{Z} = \frac{1}{N} \sum_{i=1}^N Z_i$ among agents.

The experimental results measure these metrics per agent type in order to identify the behavioral differences between agent types.

¹Our code is available at https://github.com/gliese876b/fairness_in_ssd

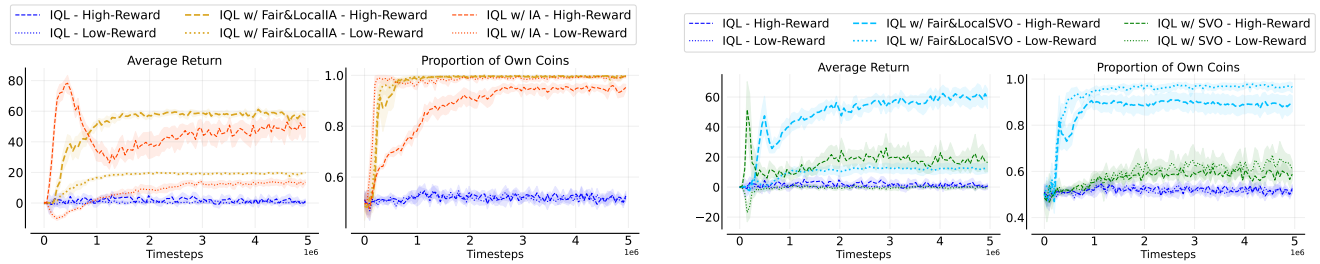


Figure 3: Performance of IA, SVO, and their Fair&Local versions in Coins with *asymmetry in coin rewards*. Fair&LocalIA and Fair&LocalSVO shows faster and more stable convergence to prosocial behavior with higher returns than their counterparts.

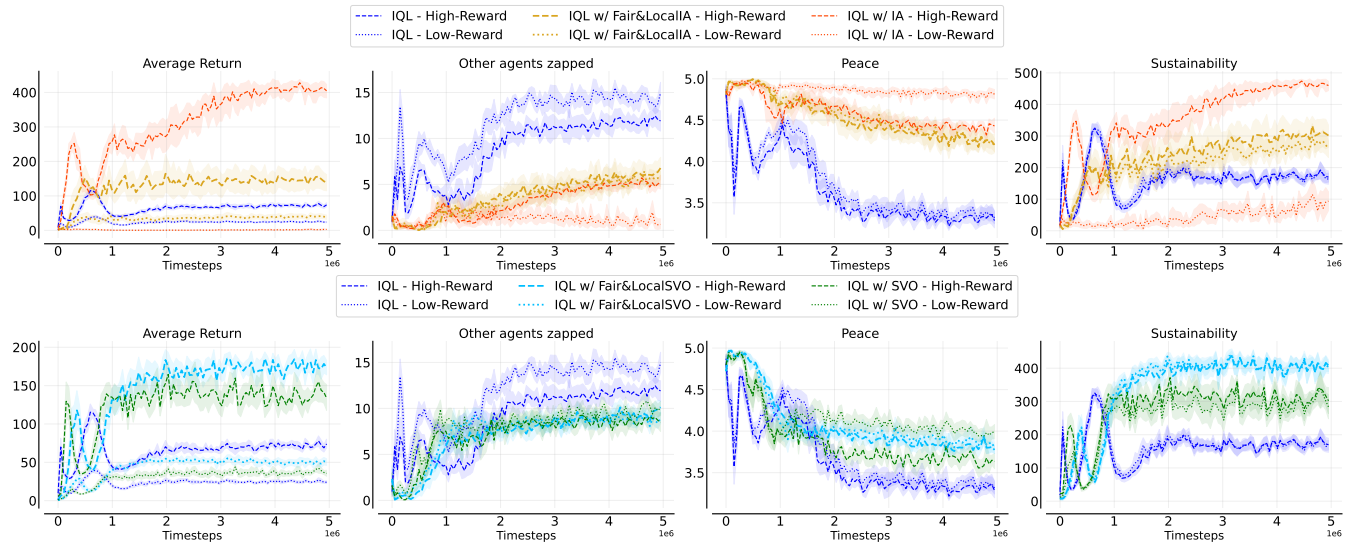


Figure 4: Performance of IA, SVO, and their Fair&Local versions in Harvest environment with *asymmetry in apple rewards*. Fair&Local versions lead to higher returns and better sustainability by eliminating the gap in peace between agent types.

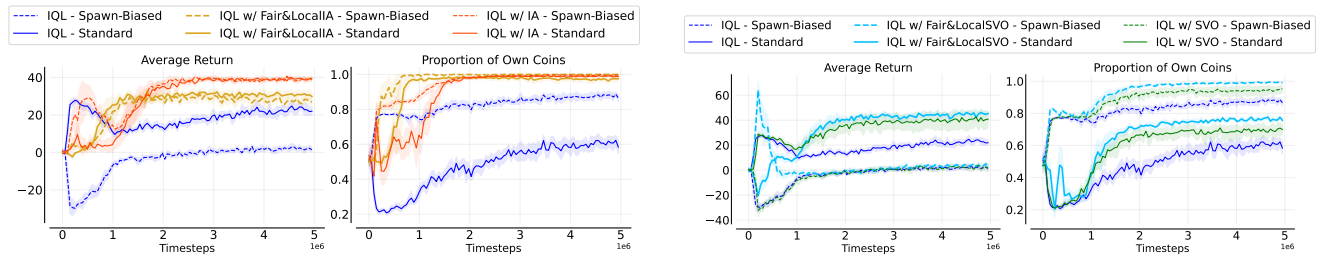


Figure 5: Performance of IA, SVO, and their Fair&Local versions in Coins environment with *asymmetry in coin spawn*. The proposed approach shows faster convergence to more cooperative policies.

5.4 Results

5.4.1 Asymmetry in Rewards. Fig. 3 shows that in Coins with one high-reward and one low-reward agent, independent learners without prosocial incentives fail to develop cooperative strategies, indiscriminately collecting coins and converging to zero average return. IA encourages some cooperation but struggles to handle the inherent asymmetry, especially for the agent with reward advantage, whereas Fair&LocalIA enables both agent types to fully cooperate,

collecting only their own coins and achieving faster convergence. SVO offers only limited improvement over the baseline IQL, while Fair&LocalSVO more effectively promotes prosocial behavior, yielding higher returns as agents primarily collect their respective coins.

In Fig. 4, the tragedy of commons is visible in Harvest with IQL, converging to mutual defection with a decrease in peace metric. It shows that the low-reward agents are more aggressive in terms of zapping. IQL with IA, gives a high punishment for low-reward agents as they fall behind of high-reward agents, causing them to

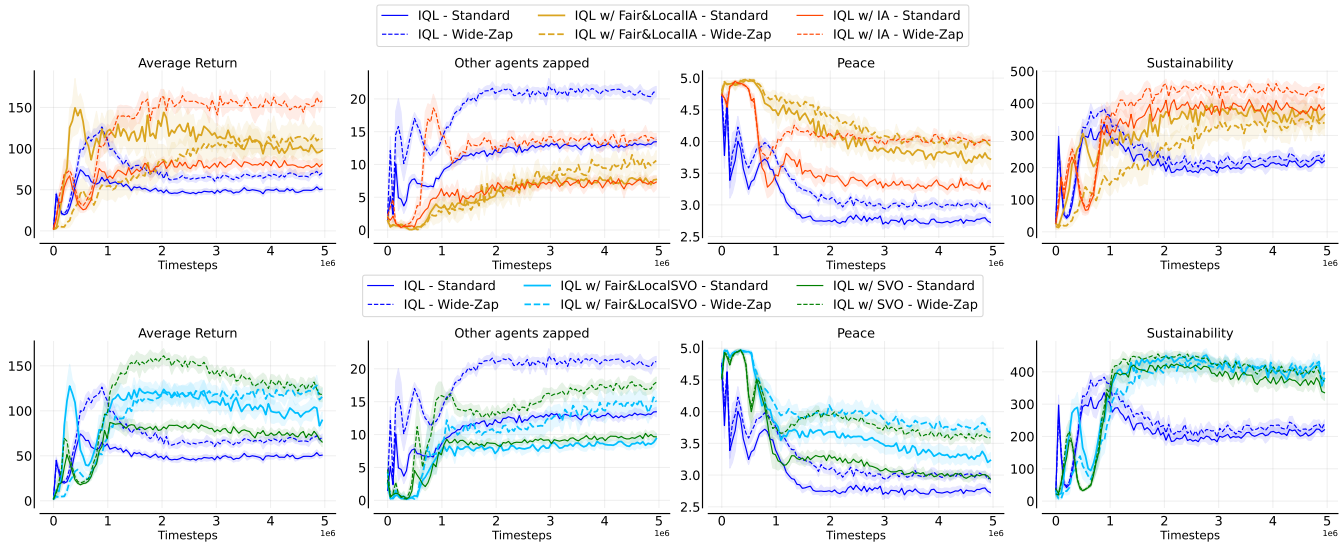


Figure 6: Performance of IA, SVO, and their Fair&Local versions in Harvest with *asymmetry in zapping*. Fair&Local versions, lead to more cooperative policies with higher peace and similar returns for each agent type.

stop moving altogether. As a result, high-reward agents collect all remaining apples, showing high sustainability among them and low peace since low-reward agents are not in the apple collection area. IQL with Fair&LocalIA, allows a better adjustment to their respective reward scales, making both types of agents actively cooperate in the environment, showing higher peace with lower zapping behavior and balanced sustainability compared to IQL with IA. While SVO shows higher returns and peace compared to IA, its wrong incentive for low-reward agents is clear in peace metric. Both types of agents zap others at the same amounts but the gap in the peace metric suggests that it is a targeted action. High-reward agents have a lower peace value, meaning that low-reward agents specifically zap the high-reward ones since this reduces the gap in rewards according to Eq. 4. This wrong incentive to defect for low-reward agents is eliminated with Fair&LocalSVO, removing the gap in peace. IQL with Fair&LocalSVO also outperforms IQL with SVO in terms of sustainability metric.

5.4.2 Asymmetry in Actions. In Coins with *asymmetry in coin spawn*, Fig. 5 shows that both agents without any social incentives act selfishly, yet this selfish act causes spawn-biased agent to collect less returns than the standard one since its defective action leads to its own coins to spawn for several time steps, and standard agent collects mostly the coins of the spawn-biased during this time period. IQL with IA shows a better result to IQL yet Fair&LocalIA converges to fully cooperative behaviors much faster than its original. Similarly, IQL with Fair&LocalSVO leads to more cooperative policies compared to IQL with SVO in this setting.

Fig. 6 shows that, in Harvest with *asymmetry in zapping*, IQL again converges to a defective setting with low peace where wide-zap agents are zapping more effectively. While showing higher returns, higher peace and higher sustainability compared to IQL, IQL with IA amplifies the gap between the wide-zap and standard agents in terms of these metrics. On the other hand, IQL w/ Fair&Local removes the gap with much less defective policies leading to higher

peace. A similar result is seen for SVO and Fair&LocalSVO, where our proposed version makes the wide-zap and standard agents perform more similarly in terms of cooperation.

Additional results analyzing the convergence of the proposed method’s key features are presented in Appendix C.

6 DISCUSSION & FUTURE WORK

In this paper, we introduced asymmetry into sequential social dilemmas, by creating differences between agents in terms of reward functions and action effects, a crucial setting for understanding how cooperation can emerge when individual and group interests conflict. Our results showed that asymmetry naturally leads to distinct behaviors across different agent types. While fairness-based methods such as IA and SVO modify agents’ reward functions to encourage prosocial behavior, in our experiments, they failed to fully address the inherent asymmetry, sometimes creating incentives to defect. In addition, these methods require constant access to other agents’ returns during training.

We proposed modifications to IA and SVO, including normalization for agent potential, agent-based weighting to account for asymmetry, and local social feedback. The resulting Fair&LocalIA and Fair&LocalSVO methods handle asymmetry more effectively, promoting higher levels of cooperation across all agent types and avoiding defection, while relying only on local information.

Future work could explore other forms or scales of asymmetry, such as differential access to resources, varied learning capabilities or larger continuous spectrum. Additionally, agent-based weights could be learned automatically during training, eliminating the need to specify them beforehand.

ACKNOWLEDGMENTS

This research is financially supported by the Scientific and Technological Research Council of Turkey through BIDEB 2219 International Postdoctoral Research Scholarship Program.

REFERENCES

- [1] John P. Agapiou, Alexander Sasha Vezhnevets, Edgar A. Duéñez-Guzmán, Jayd Matyas, Yiran Mao, Peter Sunehag, Raphael Köster, Udari Madhushani, Kavya Koppurapu, Ramona Comanescu, DJ Strouse, Michael Bradley Johanson, Sukhdeep Singh, Julia Haas, Igor Mordatch, Dean Mobbs, and Joel Z. Leibo. 2022. Melting Pot 2.0. *CoRR* abs/2211.13746 (2022). <https://doi.org/10.48550/ARXIV.2211.13746> arXiv:2211.13746
- [2] Stefano V. Albrecht, Filippos Christianos, and Lukas Schäfer. 2024. *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*. MIT Press. <https://www.marl-book.com>
- [3] James Andreoni and Hal Varian. 1999. Preplay contracting in the prisoners' dilemma. *Proceedings of the National Academy of Sciences* 96, 19 (1999), 10933–10938.
- [4] Martin Beckenkamp, Heike Hennig-Schmidt, and Frank P. Maier-Rigaud. 2007. Cooperation in symmetric and asymmetric prisoner's dilemma games. *MPI Collective Goods Preprint* 2006/25 (2007).
- [5] Matteo Bettini, Amanda Prorok, and Vincent Moens. 2024. BenchMARL: Benchmarking Multi-Agent Reinforcement Learning. *Journal of Machine Learning Research* 25, 217 (2024), 1–10. <http://jmlr.org/papers/v25/23-1612.html>
- [6] Gary Charness, Guillaume R Fréchet, and Cheng-Zhong Qin. 2007. Endogenous transfers in the Prisoner's Dilemma game: An experimental test of cooperation and coordination. *Games and Economic Behavior* 60, 2 (2007), 287–306.
- [7] Filippos Christianos, Georgios Papoudakis, and Stefano V. Albrecht. 2023. Pareto Actor-Critic for Equilibrium Selection in Multi-Agent Reinforcement Learning. *Trans. Mach. Learn. Res.* 2023 (2023). <https://openreview.net/forum?id=3AzqYa18ah>
- [8] Phillip JK Christoffersen, Andreas A Haupt, and Dylan Hadfield-Menell. 2023. Get It in Writing: Formal Contracts Mitigate Social Dilemmas in Multi-Agent RL. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*. 448–456.
- [9] 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).
- [10] David Earl Hostallero, Daewoo Kim, Sangwoo Moon, Kyunghwan Son, Wan Ju Kang, and Yung Yi. 2020. Inducing cooperation through reward reshaping based on peer evaluations in deep multi-agent reinforcement learning. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*. 520–528.
- [11] Edward Hughes, Joel Z. Leibo, Matthew Phillips, Karl Tuyls, Edgar Duéñez-Guzmán, Antonio García Castañeda, Iain Dunning, Tina Zhu, Kevin McKee, Raphael Köster, et al. 2018. Inequity aversion improves cooperation in intertemporal social dilemmas. *Advances in neural information processing systems* 31 (2018).
- [12] Marco A. Janssen and Nathan D. Rollins. 2012. Evolution of cooperation in asymmetric commons dilemmas. *Journal of Economic Behavior & Organization* 81, 1 (2012), 220–229.
- [13] 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.
- [14] Raphael Köster, Dylan Hadfield-Menell, Richard Everett, Laura Weidinger, Gillian K. Hadfield, and Joel Z. Leibo. 2022. Spurious normativity enhances learning of compliance and enforcement behavior in artificial agents. *Proceedings of the National Academy of Sciences* 119, 3 (2022), e2106028118.
- [15] Joel Z. Leibo, Edgar Duéñez-Guzmán, Alexander Sasha Vezhnevets, John P. Agapiou, Peter Sunehag, Raphael Köster, Jayd Matyas, Charles Beattie, Igor Mordatch, and Thore Graepel. 2021. Scalable Evaluation of Multi-Agent Reinforcement Learning with Melting Pot. *International conference on machine learning*. <https://doi.org/10.48550/arXiv.2107.06857>
- [16] Joel Z. Leibo, Vinicius Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. 2017. Multi-agent reinforcement learning in sequential social dilemmas. *arXiv preprint arXiv:1702.03037* (2017).
- [17] Adam Lerer and Alexander Peysakhovich. 2017. Maintaining cooperation in complex social dilemmas using deep reinforcement learning. *arXiv preprint arXiv:1707.01068* (2017).
- [18] Carlos R. Lopez, Traci M. Cihon, Aécio de Borja Vasconcelos Neto, and April Becker. 2022. An exploration of cooperation during an asymmetric iterated prisoner's dilemma game. *Behavior and Social Issues* 31, 1 (2022), 106–132.
- [19] Andrei Lupu and Doina Precup. 2020. Gifting in multi-agent reinforcement learning. In *Proceedings of the 19th International Conference on autonomous agents and multiagent systems*. 789–797.
- [20] Kevin R. McKee, Ian Gemp, Brian McWilliams, Edgar A. Duéñez-Guzmán, Edward Hughes, and Joel Z. Leibo. 2020. Social Diversity and Social Preferences in Mixed-Motive Reinforcement Learning. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '20, Auckland, New Zealand, May 9-13, 2020*, Amal El Fallah Seghrouchni, Gita Sukthankar, Bo An, and Neil Yorke-Smith (Eds.). International Foundation for Autonomous Agents and Multiagent Systems, 869–877. <https://doi.org/10.5555/3398761.3398863>
- [21] Georgios Papoudakis, Filippos Christianos, Arrasy Rahman, and Stefano V. Albrecht. 2019. Dealing with Non-Stationarity in Multi-Agent Deep Reinforcement Learning. arXiv:1906.04737 [cs.LG]
- [22] Georgios Papoudakis, Filippos Christianos, Lukas Schäfer, and Stefano V. Albrecht. 2021. Benchmarking Multi-Agent Deep Reinforcement Learning Algorithms in Cooperative Tasks. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS)*. <http://arxiv.org/abs/2006.07869>
- [23] Julien Pérolat, Joel Z. Leibo, Vinicius Flores Zambaldi, Charles Beattie, Karl Tuyls, and Thore Graepel. 2017. A multi-agent reinforcement learning model of common-pool resource appropriation. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (Eds.). 3643–3652. <https://proceedings.neurips.cc/paper/2017/hash/2b0f658cbffd284984fb11d90254081f-Abstract.html>
- [24] Alexander Peysakhovich and Adam Lerer. 2018. Prosocial Learning Agents Solve Generalized Stag Hunts Better than Selfish Ones. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*, Elisabeth André, Sven Koenig, Mehdi Dastani, and Gita Sukthankar (Eds.). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, USA / ACM, 2043–2044. <http://dl.acm.org/citation.cfm?id=3238065>
- [25] Stefan Roesch, Stefanos Leonardos, and Yali Du. 2024. Selfishness Level Induces Cooperation in Sequential Social Dilemmas. In *The 23rd International Conference on Autonomous Agents and Multi-Agent Systems: AAMAS 2024*.
- [26] Lukas Schäfer, Filippos Christianos, Josiah P. Hanna, and Stefano V. Albrecht. 2022. Decoupled Reinforcement Learning to Stabilise Intrinsically-Motivated Exploration. In *21st International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2022, Auckland, New Zealand, May 9-13, 2022*, Piotr Faliszewski, Viviana Mascardi, Catherine Pelachaud, and Matthew E. Taylor (Eds.). International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), 1146–1154. <https://doi.org/10.5555/3535850.3535978>
- [27] Thomas C. Schelling. 1973. Hockey helmets, concealed weapons, and daylight saving: A study of binary choices with externalities. *Journal of Conflict resolution* 17, 3 (1973), 381–428.
- [28] John P. Sheposh and Philip S. Gallo Jr. 1973. Asymmetry of payoff structure and cooperative behavior in the Prisoner's Dilemma game. *Journal of Conflict Resolution* 17, 2 (1973), 321–333.
- [29] Martin Smit and Fernando P. Santos. 2024. Learning Fair Cooperation in Mixed-Motive Games with Indirect Reciprocity. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI 2024, Jeju, South Korea, August 3-9, 2024*. ijcai.org, 220–228. <https://www.ijcai.org/proceedings/2024/25>
- [30] Elizaveta Tennant, Stephen Hailes, and Mirco Musolesi. 2023. Modeling Moral Choices in Social Dilemmas with Multi-Agent Reinforcement Learning. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, 19th-25th August 2023, Macao, SAR, China*. ijcai.org, 317–325. <https://doi.org/10.24963/IJCAI2023/36>
- [31] Albert W. Tucker. 1950. A two-person dilemma. *Prisoner's Dilemma* (1950).
- [32] 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.
- [33] Yicheng Wang, Xin Wang, Deyi Ren, Yaohui Ma, and Chao Wang. 2021. Effect of asymmetry on cooperation in spatial evolution. *Physical Review E* 103, 3 (2021), 032414.
- [34] Jiachen Yang, Ang Li, Mehrdad Farajtabar, Peter Sunehag, Edward Hughes, and Hongyuan Zha. 2020. Learning to incentivize other learning agents. *Advances in Neural Information Processing Systems* 33 (2020), 15208–15219.
- [35] Lei Zhang, Yan Jin, Lin Xia, Bibo Xu, and Syed Mohamad Syed Abdullah. 2022. The effects of social distance and asymmetric reward and punishment on individual cooperative behavior in dilemma situations. *Frontiers in Psychology* 13 (2022), 816168.