

Altruism and Fair Objective in Mixed-Motive Markov Games

Franck Yao-hua Xu
Orange Labs Lannion
Lannion, France
franck.xu@orange.com

Arnaud Braud
Orange Labs Lannion
Lannion, France
arnaud.braud@orange.com

Tayeb Lemlouma
IRISA
Lannion, France
tayeb.lemlouma@irisa.fr

Jean-Marie Bonnin
IRISA
Rennes, France
jean-marie.bonnin@irisa.fr

ABSTRACT

Cooperation is fundamental for society’s viability, as it enables the emergence of structure within heterogeneous groups that seek collective well-being. However, individuals are inclined to defect in order to benefit from the group’s cooperation without contributing the associated costs, thus leading to unfair situations. In game theory, social dilemmas entail this dichotomy between individual interest and collective outcome. The most dominant approach to multi-agent cooperation is the utilitarian welfare which can produce efficient highly inequitable outcomes.

This paper proposes a novel framework to foster fairer cooperation by replacing the standard utilitarian objective with Proportional Fairness. We introduce a fair altruistic utility for each agent, defined on the individual log-payoff space and derive the analytical conditions required to ensure cooperation in classic social dilemmas. We then extend this framework to sequential settings by defining a Fair Markov Game and deriving novel fair Actor-Critic algorithms to learn fair policies. Finally, we evaluate our method in various social dilemma environments.

KEYWORDS

Social Dilemma; Game Theory; Markov Game; Multi-agent Reinforcement Learning; Reinforcement Learning; Fairness

ACM Reference Format:

Franck Yao-hua Xu, Tayeb Lemlouma, Arnaud Braud, and Jean-Marie Bonnin. 2026. Altruism and Fair Objective in Mixed-Motive Markov Games. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/MPMP3285>

1 INTRODUCTION

Cooperation is a fundamental mechanism that underpins the stability and success of nearly all multi-agent systems, from biological ecosystems to complex human societies and artificial intelligence. It is the process by which independent agents coordinate their actions to pursue shared objectives, thus facilitating the emergence of complex structures and achieving a collective well-being that would

be unreachable for any single agent acting alone. The capacity for effective cooperation is not merely advantageous; it is often crucial for survival, advancement, and the efficient use of resources.

Despite its critical importance, sustained cooperation is notoriously difficult to achieve. The primary obstacle is the inherent tension between individual self-interest and the collective good—a conflict formally studied in game theory as a social dilemma [7]. In these scenarios, rational agents are frequently incentivized to defect from the group effort to pursue personal gain. This individual rationality can paradoxically culminate in collective failure, a phenomenon often described as a *tragedy of the commons*, where the pursuit of self-interest leads to profoundly unfair and suboptimal outcomes for everyone [12].

Conceptually, the opposite to the selfish defection that drives these dilemmas is altruism—the principle of acting for the benefit of others. In the field of multi-agent learning [3, 26], this altruistic motivation is commonly formalized through a utilitarian social welfare objective, which seeks to maximize the (weighted) sum of all agents’ utilities [20, 23, 25]. By focusing on the greatest total good, this approach encourages agents to undertake actions that, while not always optimal for themselves, produce the most efficient outcomes for the group as a whole, thereby offering a direct method for resolving the dilemma.

However, while the utilitarian approach can lead to efficient solutions, its exclusive focus on total welfare often ignores how rewards and sacrifices are distributed, frequently resulting in highly inequitable outcomes. For instance, a strict utilitarian calculus might conclude that a healthy individual should donate a kidney to a stranger, as the recipient’s immense gain (a saved life) outweighs the donor’s loss. Although total welfare is maximized, this outcome imposes an extreme and intuitively unfair sacrifice on one person. This highlights how pure, uncalibrated altruism can conflict with the principles of fairness necessary for stable cooperation. True, sustainable cooperation, therefore, requires more than just promoting the collective good; it demands a mechanism to ensure that the costs and benefits of working together are distributed fairly among all participants.

Fairness was first studied in resource allocation problems [4, 5, 17]. In Multi-Agent Reinforcement Learning (MARL), fairness is merely emerging as an active research area and is usually treated alongside cooperation. The popular approaches to promote fairness in MARL use reward shaping [2, 8, 10, 14] or role-based reward shaping [19]. However, these methods are often environment-dependent, reward instantaneous fair behavior, and neglect the fairness on the



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/MPMP3285>

full trajectory. While related works [15, 21] address similar problems, these studies are limited to finite trajectories and lack a deep analysis of the policy gradient.

To address these limitations and the need for a fairer approach, this paper introduces a novel framework designed to explicitly foster more equitable cooperation with heterogeneous agents. We move beyond simple reward maximization and propose a new objective based on Proportional Fairness, a concept that effectively balances efficiency with equity.

Our contributions are threefold. First, we introduce a fair altruistic utility for each agent, defined on the individual log-payoff space, and derive the analytical conditions needed to ensure cooperation in classic social dilemmas. Second, we extend this framework to sequential settings by defining a Fair Altruistic Markov Game with infinite horizon. Finally, based on this new game formulation, we derive novel fair Actor-Critic algorithms designed to learn policies that achieve both high collective rewards and equitable outcomes, and we validate our method in various social dilemma environments.

2 PRELIMINARY

2.1 Notation

We denote the set of real numbers by \mathbb{R} and positive real numbers by \mathbb{R}_+ . We denote d -dimensional real-numbered vector spaces as \mathbb{R}^d . By $\mathcal{P}(A)$, we denote the power set of a set A . For any probability measure \mathbb{P} , we denote the probability of an event $X = x$ as $\mathbb{P}(X = x)$. Similarly, we write $\mathbb{P}(X = x | Y = y)$ for conditional probabilities. When it is clear, which random variable is referred to, we regularly omit it to shorten notation, i.e. $\mathbb{P}(X = x) = \mathbb{P}(x)$.

We denote that a random variable X follows a probability distribution p by $X \sim p$. For any random variable $X \sim p$, we denote by $\mathbb{E}_{X \sim p}[X]$ its expectation. We denote the set of probability distributions over some measurable space \mathcal{A} as $\Delta(\mathcal{A})$. We write $|A|$ for the cardinality of a finite set A . We use the vector notation \vec{a} to refer to any element in a cartesian product, i.e. for any sets A_1, \dots, A_n , $\vec{a} = (a_1, \dots, a_n) \in \times_{i=1}^n A_i$.

2.2 Normal-form games

We define a normal-form game G by the tuple $\langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$, where $N = \{1, 2, \dots, n\}$ is the set of players, S_i is the set of pure strategies available to player i . Let $\mathcal{G}_{N,S}$ be the set of all normal form games with a fixed set of players $N = \{1, \dots, n\}$ and fixed sets of pure strategies $S = (S_i)_{i \in N}$. $p_i : S \rightarrow \mathbb{R}$ is the payoff function for player i , which assigns a real-valued utility to each strategy profile $s \in S$. In this work, we assume the utility to be bounded.

Nash equilibrium. In game theory, a fundamental solution concept is the Nash equilibrium where no player can benefit by changing their strategy while the other players keep theirs unchanged. Formally, a strategy profile $s^* = (s_1^*, s_2^*, \dots, s_n^*)$ is a Nash equilibrium if, for every player $i \in N$ and every alternative strategy $s_i \in S_i$:

$$p_i(s_i^*, s_{-i}^*) \geq p_i(s_i, s_{-i}^*).$$

2.3 Altruistic games

Altruism in game theory is often framed through social welfare functions (SWF), which aggregate individual utilities into collective wealth. While utilitarian SWFs maximize overall efficiency, they can permit extreme inequality. In contrast, "fair" SWFs satisfy axioms of efficiency (Pareto improvement), equity (Pigou-Dalton), and anonymity to ensure equitable distribution [22].

Following intuition, it is the smallest proportion of social welfare that needs to be offered to each player to make the socially optimal outcome a stable Nash equilibrium. In other terms, it is a measure of how much incentive is needed to make a selfish person act cooperatively. It's a way to measure the tension between an individual's self-interest and the collective good. This is the α -selfish game, the most studied altruistic game [23].

In this work, we adopt a slightly different definition [6].

Definition 2.1. (α -altruistic extension of a Normal-form Game) Given any normal-form game $G = \langle N, \{S_i\}_{i \in N}, \{p_i\}_{i \in N} \rangle$, we can induce an altruistic game $G(\alpha) \doteq \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ where, $u_i(s) \doteq (1 - \alpha)p_i(s) + \alpha SW(s)$. The altruism level of a strategic game G is then defined as:

$$\alpha_G = \inf_{\alpha} \{ \alpha \in \mathbb{R}_+ | G(\alpha) \text{ is } \alpha\text{-altruistic} \} \quad (1)$$

where, G is α -altruistic if, for some $\alpha \geq 0$, a pure Nash equilibrium of $G(\alpha)$ is a social optimum of G .

2.4 Social dilemmas

Social dilemmas are a class of normal-form games where players create a collectively worse outcome for the entire group when acting in their own self-interest. It highlights the dualism between individual preferences and the collective good.

More formally, social dilemmas are symmetrical two-players normal-form games with payoff matrices similar as described in Table 1.

Table 1: Outcome categories in the payoff matrix

	C	D
C	R, R	S, T
D	T, S	P, P

where R is the reward for mutual cooperation, P the Punishment for mutual defection, S the sucker's payoff for cooperating while the other defects and T the temptation to defect while the other cooperates.

A situation is defined as a social dilemma if the payoffs satisfy the following preference inequalities.

2.5 Proportional Fairness

The resource allocation problem is the problem of allocating limited resources to competing tasks or individuals while maximizing an objective function – for instances, performance, equity, or costs. The Proportional fairness (PF) objective was introduced to provide a compromise between efficiency and some kind of fairness, and was firstly introduced in telecommunication for rate control [4, 9, 17].

Formal Definition: Let N be the set of agents, $X \subseteq \mathbb{R}_+^N$ the set of all feasible allocations. Let us denote by $u_i : X \rightarrow \mathbb{R}_+$ the utility

Table 2: Set of inequalities that define social dilemmas.

Inequality	Preference
$R > P$	The individual prefers mutual cooperation (C, C) to mutual defection (D, D).
$R > S$	The individual prefers mutual cooperation to unilateral cooperation (C, D).
$2R > T + S$	The group prefers mutual cooperation to unilateral cooperation/defection.
$T > R$ or $P > S$	The individual prefers unilateral defection (D, C) to mutual cooperation (aka, greed) or the individual prefers mutual defection to unilateral cooperation (aka fear).

of agent i . A solution $x^* \in X$ is proportionally fair if for any $x \in X$, the sum of proportional variation is negative:

$$\sum_{i=1}^N \frac{u_i(x) - u_i(x^*)}{u_i(x^*)} \leq 0$$

Proportional Fairness has been widely studied in the literature. In particular, it has been proven that a proportional fair allocation x^* is a solution of the following optimization problem:

$$\max_{x \in X} \sum_{i=1}^N \log u_i(x)$$

Proportional fairness is directly linked to the Nash welfare function [5, 16], which seeks to maximize the product of agents’ utilities, as maximizing the product of utilities is equivalent to maximize the sum of their logarithms in some sense.

2.6 Markov game

For modeling sequential decision-making in multi-agent settings, the standard framework is the Markov Game (also known as a Stochastic Game).

A Markov game \mathcal{M} is defined by the following tuple: $(N, \mathcal{S}, \{\mathcal{A}_i\}_{i \in N}, P, \rho_0, \{r_i\}_{i \in N}, \gamma)$, with N the finite set of n players, \mathcal{S} the space of all possible states of the environment, $\{\mathcal{A}_i\}_{i \in N}$ the individual action set for each player, $P(s', \vec{a}, s) = \mathbb{P}(s' | s, \vec{a})$ the transition probability kernel, $\rho_0 \in \Delta(\mathcal{S})$ the initial state distribution, $\{r_i\}_{i \in N}$ the individual reward for each player and $\gamma \in [0, 1)$ the discount factor. In this work, we assume the state and action spaces to be finite and the reward functions to be positive and bounded by R_{\max} .

3 METHODOLOGY

In this section, we extend the α -altruistic Game by rescaling the payoff function to use the Proportional Fairness function. Then, we study the effect of such a transformation on simple games, specifically two-player social dilemmas.

3.1 Redefining Altruistic Games

The former definition of the α -altruistic game is not directly compatible with the proportional fairness function because summing individual rewards with log-rewards is not a homogeneous operation. To solve this, we apply a function F_i to selfish rewards so

they operate on the same scale as the social welfare function. For instance, we use $F_i = \log$ for PF. Additionally, the logarithm function restricts the considered values to strictly positive real values. Hence our alternative definition,

Definition 3.1. (α -altruistic extension of a Normal-form Game) Given any normal-form game $G = \langle N, \{S_i\}_{i \in N}, \{p_i\}_{i \in N} \rangle$, we can induce a fair altruistic game $G(\alpha) \doteq \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ where, $u_i(s) \doteq (1 - \alpha)F_i(p_i(s) - m_p) + \alpha SW(s)$, where $m_p = \inf_{i,s} [u_i(s)]$ is the min payoff a player can get over all possible strategy (which is assumed to exist) and $F_i : \mathbb{R} \rightarrow \mathbb{R}$ is an increasing function applied on player’s payoff. The altruism level of a strategic game G is then defined as:

$$\alpha_G = \inf_{\alpha} \{ \alpha \in \mathbb{R}_+ | G \text{ is } \alpha\text{-altruism} \} \quad (2)$$

where, G is α -altruism if, for some $\alpha \geq 0$, a pure Nash equilibrium of $G(\alpha)$ is a social optimum of G .

In the rest of this work, we set $F_i = \log$ for all players i .

3.2 Analysis of Social Dilemmas

Let $G \in \mathcal{G}_{N,S}$ be a normal form game where N is the set of players, S is the set of strategy profiles and R is the payoff function, where $R_i(s)$ is the payoff to player i . We assume all payoffs are strictly positive, i.e., $R_i(s) > 0$ for all i, s .

We define a transformed game $G'(\alpha) = \langle N, S, R' \rangle$, where the transformed payoff function R' is defined as:

$$R'_i = (1 - \alpha) \log R_i + \alpha \sum_{j \in N} \log R_j = \log R_i + \alpha \sum_{j \neq i} \log R_j$$

We now apply this framework to derive the threshold for α in the Prisoner’s Dilemma, the Stag Hunt, and the Game of Chicken.

For the case of 2-player games ($N = 2$), the rule for Player 1 and Player 2 simplifies to

$$R'_1 = \log R_1 + \alpha \log R_2,$$

$$R'_2 = \log R_2 + \alpha \log R_1.$$

The parameter α represents the degree of payoff interdependence. If $\alpha = 0$, the game is unchanged ($R'_i = R_i$). As α increases, a player’s transformed payoff becomes increasingly influenced by the other player’s original payoff.

Our analysis will identify the conditions on α for the mutual cooperation outcome (C, C) to be an equilibrium in $G'(\alpha)$. For this to be true, no player must have a profitable unilateral deviation. For a symmetric 2-player game, we only need to ensure that Player 1 does not gain by defecting, assuming Player 2 cooperates. This condition is

$$R'_1(C, C) \geq R'_1(D, C).$$

THEOREM 3.2. *The altruism level of social dilemma is*

$$\alpha_G = \begin{cases} 0 & \text{if } T \leq R, \\ \frac{\log T - \log R}{\log R - \log S} & \text{if } T > R. \end{cases} \quad (3)$$

PROOF. In the transformed game, the payoff becomes

$$\begin{aligned} R'_2(C, C) &= R'_1(C, C) = (1 + \alpha) \log R \\ R'_1(D, C) &= \log T + \alpha \log S \\ R'_2(D, C) &= \log S + \alpha \log T \end{aligned}$$

To make (C, C) a Nash Equilibrium, the condition $R'_1(C, C) \geq R'_1(D, C)$ must hold.

The Prisoner’s Dilemma. The Prisoner’s Dilemma is defined by the preference order $T > R > P > S$. The dilemma is that mutual defection (D, D) is the unique Nash Equilibrium, despite mutual cooperation (C, C) being a Pareto-superior outcome.

The condition becomes $(1 + \alpha) \log R \geq \log T + \alpha \log S$. Solving this inequality for α yields to the following condition

$$\alpha > \frac{\log T - \log R}{\log R - \log S}$$

The Stag Hunt. The Stag Hunt is defined by the preference order $R > P > S$ and $T \leq R$. This game has two pure-strategy equilibrium: (C, C) and (D, D) . The cooperative outcome is already an equilibrium. The analysis thus investigates whether the transformation preserves this equilibrium.

The condition is identical to the one above: $R'_1(C, C) \geq R'_1(D, C)$. The resulting inequality is also $(1 + \alpha) \log R \geq \log T + \alpha \log S$. However, the implications are different. In a Stag Hunt, $T \leq R$, so $\log T - \log R$ is negative, while $\log R - \log S$ is positive. The ratio $\frac{\log T - \log R}{\log R - \log S}$ is therefore negative. As α is always positive, the condition is met for any $\alpha \in [0, 1]$. This confirms that the transformation never change the cooperative equilibrium.

The Game of Chicken. The Game of Chicken is defined by the preference order $T > R > S \geq P$. Mutual cooperation (C, C) is not an equilibrium because each player is tempted to defect to achieve the highest payoff.

To make (C, C) a Nash Equilibrium, the following condition must be met

$$\alpha > \frac{\log T - \log R}{\log R - \log S}$$

Therefore, cooperation becomes a Nash Equilibrium whenever the following condition is met

$$\alpha = \begin{cases} 0 & \text{if } T \leq R, \\ \frac{\log T - \log R}{\log R - \log S} & \text{if } T > R. \end{cases}$$

To ensure consistency with the definition of altruism level $\alpha \leq 1$, we require $TS \leq R^2$, which holds for standard payoff matrices. \square

In each case where cooperation is not initially stable (Prisoner’s Dilemma, Chicken), the required threshold for α is a ratio of logarithmic payoffs. The numerator, $\log(T/R)$, represents the "temptation" to defect from the cooperative state. The denominator reflects the consequences of defection. The transformation effectively allows players to internalize the externalities of their actions, and the threshold for α quantifies the precise degree of internalization needed to make cooperation individually rational.

The use of proportional fairness combined with altruism is evident in traditional static social dilemmas. However, the simplicity of Normal-form games is a limitation. Indeed, real-world phenomena rarely occur in a single step, but rather evolve over time. To accurately model such dynamic systems, the analysis must move beyond such formalism to a framework like the Markov Game.

4 FAIR ALTRUISTIC MARKOV GAME

In this section, we propose and analyze a novel objective function integrating Proportional Fairness and Altruism within the Markov Game framework and derive policy gradient theorems¹.

4.1 Value functions

To extend our approach to Markov Games, we must establish a dynamic analogy to the players’ static payoff function for a given policy profile $\vec{\pi}$. This can be achieved using the individual state-value function which represents the expected cumulated discounted rewards for a single player over all possible infinite trajectories.

One challenge in defining the state-value function is to construct a single, consistent probability measure, $\mathbb{P}_{\vec{\pi}}$, over the entire infinite set of trajectories, $\mathcal{T} = (\mathcal{S} \times \mathcal{A})^\infty$. We do this by defining the probability of any finite sequence and then extending it to the infinite case using the Ionescu-Tulcea Theorem.

First of all, let define the discounted return R_i for each agent $i \in N$ and for a given trajectory $\tau = (s_0, \vec{a}_0, s_1, \vec{a}_1, \dots) \in \mathcal{T}$

$$R_i(\tau) = \sum_{t=0}^{\infty} \gamma^t r_i(s_t, \vec{a}_t).$$

As the reward function is bounded, the discounted return converges absolutely.

The state-value function of agent i is the expected return with respect to $\mathbb{P}_{\vec{\pi}}$ given the initial state $s_0 \in \mathcal{S}$,

$$V_i^{\vec{\pi}}(s) = \mathbb{E}_{\tau \sim \mathbb{P}_{\vec{\pi}}} [R_i(\tau) \mid s_0 = s].$$

Similarly, the action-value function of agent i is defined as the expected discounted return given the initial state and the first joint action,

$$Q_i^{\vec{\pi}}(s, \vec{a}) = \mathbb{E}_{\tau \sim \mathbb{P}_{\vec{\pi}}} [R_i(\tau) \mid s_0 = s, \vec{a}_0 = \vec{a}]$$

In a Markov Game, the Bellman operator is a fundamental concept used in multi-agent reinforcement learning to define the value of a state for a specific agent.

Definition 4.1. In a Markov Game with a set of agents N and joint stationary policies $\vec{\pi}$, the Bellman operator for agent i , denoted as $T_i^{\vec{\pi}}$, is applied to any state-value function $V_i(s)$ and is defined as:

$$(T_i^{\vec{\pi}} V_i)(s) = \sum_{\vec{a} \in \mathcal{A}} \pi(\vec{a} \mid s) \left(r_i(s, \vec{a}) + \gamma \sum_{s' \in \mathcal{S}} \mathbb{P}(s' \mid s, \vec{a}) V_i(s') \right).$$

We can then write the Bellmann theorem which gives a key equation verified by $V_i^{\vec{\pi}}$

THEOREM 4.2. *For any joint policy $\vec{\pi}$, and any state-value function associated with $\vec{\pi}$, the following equation called Bellman equation holds*

$$T_i^{\vec{\pi}} V_i^{\vec{\pi}} = V_i^{\vec{\pi}}$$

4.2 Fair objective

From the traditional Markov Game framework, we can induce a new objective that balances fairness and efficiency and extends the altruistic Normal-form game to the Markov game.

¹More detailed proofs in Appendix on arxiv <https://arxiv.org/abs/2602.08389>.

Definition 4.3. For each agent i , we define the proportional fair state value function :

$$V^{\vec{\pi}, \text{Prop}}(s) = \sum_{j=1}^n \log V_j^{\vec{\pi}}(s).$$

The state value functions play the role of the payoff function and the proportional fairness function is applied on the state value function. Therefore, we can define a new objective for each agent depending on the altruism level:

Definition 4.4. The objective for each agent is to maximize

$$\begin{aligned} J_i(\vec{\pi}) &= \mathbb{E}_{s_0 \sim \rho_0} \left[(1 - \alpha) \log V_i^{\vec{\pi}}(s_0) + \alpha V^{\vec{\pi}, \text{Prop}}(s_0) \right] \\ &= \mathbb{E}_{s_0 \sim \rho_0} \left[\sum_{j=1}^N c_i(j) \log V_j^{\vec{\pi}}(s_0) \right] \end{aligned}$$

where

$$c_i(j) = \begin{cases} 1 & \text{if } j = i, \\ \alpha & \text{if } j \neq i. \end{cases}$$

4.3 Policy gradient

Policy gradient methods are a class of reinforcement learning algorithms that directly learn and optimize a parameterized policy, $\pi_\theta(a | s)$, without consulting a value function. The core principle is to adjust the policy's parameters, θ , to increase the probability of actions that yield high cumulative rewards and decrease the probability of actions that result in low rewards.

Let us suppose that agent i 's policy π_i is parameterized by θ_i such that π_{θ_i} is differentiable with respect to θ_i .

Therefore, the goal of agent i is to maximize the objective $J_i(\vec{\theta}) = \mathbb{E}_{s_0 \sim \rho_0} [\sum_{j=1}^N c_i(j) \log V_j^{\vec{\theta}}(s_0)]$. We can demonstrate the differentiability of this objective function with respect to any parameter θ as summarized in the following theorem.

THEOREM 4.5. (Policy gradient theorem.) Let \mathcal{M} be a Markov game and $J_i(\vec{\theta}) = \mathbb{E}_{s_0 \sim \rho_0} [\sum_{j=1}^N c_i(j) \log V_j^{\vec{\theta}}(s_0)]$ the objective defined previously and for all agent i , let us suppose $Q_i^{\vec{\pi}}$ is differentiable with respect to θ_i under mild conditions. The gradient of J_i with respect to θ_i is given by

$$\nabla_{\theta_i} J_i(\vec{\theta}) = \mathbb{E}_{\tau \sim \mathbb{P}^{\vec{\theta}}} \left[\sum_{t=0}^{\infty} \sum_{j=1}^N c_i(j) \frac{Q_j^{\vec{\theta}}(s_t, \vec{a}_t)}{V_j^{\vec{\theta}}(s_0)} \nabla_{\theta_i} \log \pi_{\theta_i}(a_{i,t} | s_t) \right]$$

PROOF. Let us derive the objective for agent i using the chain rule,

$$\nabla_{\theta_i} J_i(\vec{\theta}) = \mathbb{E}_{s_0 \sim \rho_0} \left[\sum_{j=1}^N c_i(j) \frac{\nabla_{\theta_i} V_j^{\vec{\theta}}(s_0)}{V_j^{\vec{\theta}}(s_0)} \right]$$

Then, we need to find the gradient $\nabla_{\theta_i} V_j^{\vec{\theta}}(s_0)$. To do so, we start from the Bellman equation, differentiate it and define a contraction mapping on the state value function gradient space and apply the Banach-Picard Fixed-Point Theorem to get the formula of the gradient.

1. Differentiate the Bellman equation. r_i, P , and policies π_{θ_j} for $j \neq i$ do not depend on θ_i . Then, the Bellman equation can be differentiated using the product rule into a recursive equation:

$$\begin{aligned} \nabla_{\theta_i} V_j^{\vec{\theta}}(s) &= \sum_{\vec{a} \in \mathcal{A}} \nabla_{\theta_i} \pi_{\vec{\theta}}(\vec{a} | s) \cdot \left(r_j(s, \vec{a}) + \gamma \sum_{s' \in \mathcal{S}} \mathbb{P}(s' | s, \vec{a}) V_j^{\vec{\theta}}(s') \right) \\ &\quad + \pi_{\vec{\theta}}(\vec{a} | s) \cdot \gamma \sum_{s' \in \mathcal{S}} \mathbb{P}(s' | s, \vec{a}) \nabla_{\theta_i} V_j^{\vec{\theta}}(s') \end{aligned}$$

Let us define a function $G_{i,j}^{\vec{\theta}}(s)$ which represents the immediate gradient component at state s

$$\begin{aligned} G_{i,j}^{\vec{\theta}}(s) &:= \sum_{\vec{a} \in \mathcal{A}} \pi_{\vec{\theta}}(\vec{a} | s) (\nabla_{\theta_i} \log \pi_{\theta_i}(a_i | s)) Q_j^{\vec{\theta}}(s, \vec{a}) \\ &= \mathbb{E}_{\vec{a} \sim \pi_{\vec{\theta}}} [\nabla_{\theta_i} \log \pi_{\theta_i}(a_i | s) Q_j^{\vec{\theta}}(s, \vec{a})] \end{aligned}$$

With this, our recursive equation can be simplified into

$$\nabla_{\theta_i} V_j^{\vec{\theta}}(s) = G_{i,j}^{\vec{\theta}}(s) + \gamma \sum_{\vec{a} \in \mathcal{A}} \pi_{\vec{\theta}}(\vec{a} | s) \sum_{s'} \mathbb{P}(s' | s, \vec{a}) \nabla_{\theta_i} V_j^{\vec{\theta}}(s')$$

2. Define Bellman operator for the value function gradient

This equation defines a fixed-point relationship. Let $g_j : \mathcal{S} \rightarrow \mathbb{R}^{|\theta_j|}$ be a vector field of gradients. We can define the Bellman operator for the value function gradient, $\mathfrak{I}_{i,j}^{\vec{\theta}}$, as

$$(\mathfrak{I}_{i,j}^{\vec{\theta}} g_j)(s) := G_{i,j}^{\vec{\theta}}(s) + \gamma \mathbb{E}_{\vec{a} \sim \pi_{\vec{\theta}}(\cdot | s), s' \sim P(\cdot | s, \vec{a})} [g_j(s')]$$

The gradient we are looking for, $\nabla_{\theta_i} V_j^{\vec{\theta}}$, is a fixed point of this operator

$$\nabla_{\theta_i} V_j^{\vec{\theta}} = \mathfrak{I}_{i,j}^{\vec{\theta}}(\nabla_{\theta_i} V_j^{\vec{\theta}})$$

3. Policy Gradient Theorem for Markov Games

The previous operator $\mathfrak{I}_{i,j}^{\vec{\theta}}$ defined a contraction mapping on the state-value function gradient space. By the **Banach-Picard fixed-point theorem**, it has a unique fixed point. This proves that the gradient vector field $\nabla_{\theta_i} V_j^{\vec{\theta}}$ is unique. Moreover, the **Banach-Picard fixed-point theorem** also states that for any sequence $(g_n)_{n \in \mathbb{N}}$ where $g_{n+1} = \mathfrak{I}_{i,j}^{\vec{\theta}}(g_n)$ and $g_0 : \mathcal{S} \rightarrow \mathbb{R}^{|\theta_j|}$, $(g_n)_{n \in \mathbb{N}}$ converges $\nabla_{\theta_i} V_j^{\vec{\theta}}$.

We can find the solution by unrolling the recursion and starting with $g_0 = 0$. As g_n converges to $\nabla_{\theta_i} V_j^{\vec{\theta}}$, the following equality holds

$$\nabla_{\theta_i} V_j^{\vec{\theta}}(s_0) = g_\infty(s_0) = \mathbb{E}_{\tau \sim \mathbb{P}^{\vec{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^t G_{i,j}^{\vec{\theta}}(s_t) \mid s_0 \right]$$

Substituting the definition of $G_{i,j}^{\vec{\theta}}(s_t)$:

$$\nabla_{\theta_i} V_j^{\vec{\theta}}(s_0) = \mathbb{E}_{\tau \sim \mathbb{P}^{\vec{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^t \nabla_{\theta_i} \log \pi_{\theta_i}(a_{i,t} | s_t) Q_j^{\vec{\theta}}(s_t, \vec{a}_t) \mid s_0 \right]$$

We can simplify this equation and inject the expression of $\nabla_{\theta_i} V_j^{\vec{\theta}}(s_0)$ into the former gradient, we get the **Fair Policy Gradient Theorem**

$$\nabla_{\theta_i} J_i(\vec{\theta}) = \mathbb{E}_{\tau \sim \mathbb{P}^{\vec{\theta}}} \left[\sum_{t=0}^{\infty} \gamma^t \nabla_{\theta_i} \log \pi_{\theta_i}(a_{i,t} | s_t) \left(\sum_{j=1}^N c_i(j) \frac{Q_j^{\vec{\theta}}(s_t, \vec{a}_t)}{V_j^{\vec{\theta}}(s_0)} \right) \right].$$

□

One common technique in reinforcement learning to stabilize learning is to add a baseline to the gradient. The same technique can be applied in our case.

THEOREM 4.6. (*Fair Altruistic Advantage Policy Gradient*). *The gradient can be rewritten*

$$\nabla_{\theta_i} J_i(\vec{\theta}) = \mathbb{E}_{\tau \sim \mathbb{P}^{\vec{\theta}}} \left[\sum_{t=0}^{\infty} A_i^{F, \vec{\theta}}(s_t, \vec{a}_t, s_0) \nabla_{\theta_i} \log \pi_{\theta_i}(a_{i,t} | s_t) \right]$$

where $A_i^{F, \vec{\theta}}(s, \vec{a}, s') = \sum_{j=1}^N c_i(j) \frac{A_j^{\vec{\theta}}(s, \vec{a})}{V_j^{\vec{\theta}}(s')}$ is the fair altruistic advantage function of agent i .

PROOF. The proof relies on showing that the expected value of the gradient term associated with any baseline function of s_t is zero. Let us consider the expectation of this baseline term at a single timestep t for a given state s_t . The expectation is over the joint action \vec{a}_t drawn from the joint policy $\vec{\pi}(\cdot | s_t)$.

We want to show that for any $f : \mathcal{S} \rightarrow \mathbb{R}$,

$$\mathbb{E}_{\vec{a}_t \sim \pi_{\vec{\theta}}(\cdot | s_t)} \left[\nabla_{\theta_i} \log \pi_{\theta_i}(a_{i,t} | s_t) f(s_t) \right] = 0$$

Notice that $f(s_t)$ does not depend on \vec{a}_t so we can pull it out of the expectation,

$$f(s_t) \cdot \mathbb{E}_{\vec{a}_t \sim \pi_{\vec{\theta}}(\cdot | s_t)} \left[\nabla_{\theta_i} \log \pi_{\theta_i}(a_{i,t} | s_t) \right].$$

The inner function in the expectation depends only on the action of agent i . Then, we can remove the other actions' influence from the expectation,

$$f(s_t) \cdot \mathbb{E}_{a_{i,t} \sim \pi_{\theta_i}(\cdot | s_t)} \left[\nabla_{\theta_i} \log \pi_{\theta_i}(a_{i,t} | s_t) \right].$$

Let us expand the expectation

$$f(s_t) \sum_{a_{i,t} \in \mathcal{A}_i} \pi_{\theta_i}(a_{i,t} | s_t) \nabla_{\theta_i} \log \pi_{\theta_i}(a_{i,t} | s_t).$$

Using the log-derivative trick

$$\begin{aligned} \sum_{a_{i,t}} \pi_{\theta_i}(a_{i,t} | s_t) \nabla_{\theta_i} \log \pi_{\theta_i}(a_{i,t} | s_t) &= \sum_{a_{i,t}} \nabla_{\theta_i} \pi_{\theta_i}(a_{i,t} | s_t) \\ &= \nabla_{\theta_i} \sum_{a_{i,t}} \pi_{\theta_i}(a_{i,t} | s_t) \\ &= \nabla_{\theta_i} (1) \\ &= 0. \end{aligned}$$

Since the expected value of the baseline term is zero, subtracting it from the term inside the expectation in the original policy gradient theorem does not change the total expectation. In this work, we use the baseline $f(s_t) = \sum_{j=1}^N c_i(j) \frac{V_j^{\vec{\theta}}(s_t)}{V_j^{\vec{\theta}}(s_0)}$. We then get the advantage

form of the fair policy gradient

$$\nabla_{\theta_i} J_i(\vec{\theta}) = \mathbb{E}_{\tau \sim \mathbb{P}^{\vec{\theta}}} \left[\sum_{t=0}^{\infty} A_i^{F, \vec{\theta}}(s_t, \vec{a}_t, s_0) \nabla_{\theta_i} \log \pi_{\theta_i}(a_{i,t} | s_t) \right]$$

□

In particular, when $\alpha = 1$, all agents try to maximize the same objective

$$J_i(\vec{\theta}) = \mathbb{E}_{s_0 \sim \rho_0} \left[V^{\vec{\theta}, \text{Prop}}(s_0) \right]$$

and the policy gradient becomes

$$\nabla_{\theta_i} J_i(\vec{\theta}) = \mathbb{E}_{\tau \sim \mathbb{P}^{\vec{\theta}}} \left[\sum_{t=0}^{\infty} A_i^{F, \vec{\theta}}(s_t, \vec{a}_t, s_0) \nabla_{\theta_i} \log \pi_{\theta_i}(a_{i,t} | s_t) \right]$$

where $A_i^{F, \vec{\theta}}(s, \vec{a}, s') = \sum_{j=1}^N \frac{A_j^{\vec{\theta}}(s, \vec{a})}{V_j^{\vec{\theta}}(s')}$ is the fair advantage function.

From this policy gradient theorem, we can derive multi-agent Actor-Critic algorithms that learn fair policies. We try to validate our method by comparing measured fairness and efficiency of multi-agent systems in a well-known environment.

5 EXPERIMENTS

In this section, we present the experiments. In a nutshell, we study the effects of the fair objective on the agents' performances and the group's fairness under various altruism level in a popular mixed-motiv game called 'CleanUp'.

5.1 Environment

The 'CleanUp' environment was introduced in [13] and is now part of Melting Pot 2.0 [1]. It is a commonly used multi-agent substrate designed to explore complex social dilemmas. It is built as a resource management game where multiple agents must cooperate to maintain a renewable resource pool. In 'CleanUp', agents have to harvest as many apples as they can. Agents only gain an immediate positive reward for harvesting one apple.




Figure 1: The CleanUp environment. Agents' observation window is represented in lighter tiles.

However, the growth rate of apples depends on how clean the river is. Whenever the pollution level becomes too high, the growth rate drops to zero and no additional apples can appear. To prevent this, agents have to clean the river and reduce pollution, benefiting




Figure 2: Comparison between Prop. Fairness and Util. Welfare objective trained with MAPPO in the fully cooperative setting ($\alpha = 1$). Solid lines is obtained by averaging the metric over a rolling window of 50 runs. Shaded areas represent the min and the max of the respective metric over the same rolling window.

all players. Thus, agents have to sacrifice short-term gains for long-term rewards as they cannot harvest apples while cleaning the river at the same time. This setup forces agents to navigate a classic social dilemma, balancing the selfish incentive of individual reward maximization (harvesting) against the collective necessity of long-term sustainability (cleaning), thereby testing their capacity for cooperation and pro-social behavior.

5.2 Setup

The environment is setup with partial observability. Agents can only view a part of the world consisting of a squared window of size 11×11 pixels around their position. In this work, we instantiate any environment with 7 agents with random initial positions.

Our codebase² leverages the Jax-based 'CleanUp' environment provided by the SocialJax framework [11]. We employ decentralized Actor-Critic algorithms, where each agent is modeled by its own set of neural networks, ensuring non-shared parameters (i.e., independent learning). This non-shared parameterization allows us to simulate a broader range of strategy profiles and better capture the full complexity of the CleanUp social dilemma. In particular, we use two convolutional layers to handle the spatial nature of observations followed by two fully connected layers. For each agent, the actor network represents the parameterized policy while the critic network estimates the individual state-value function.

We ran our training algorithms over 10 parallel environments. The episode length is set to 100 timesteps. We keep the episode length quite low on purpose to study the impact of the initial position on the learning process. Players spawned close to the apple area have an initial advantage over those closer to the river even with random policies. The total timestep is set to 3×10^5 . The learning rate is set to 0.001 for PPO-like algorithms [27] and 0.01 for A2C-like algorithms, with a decay schedule applied on both to reach a final value of 0.00001. In PPO-like algorithms, we make advantage of the entropy regularization to improve exploration.

In our experiment, we study the performance of the following algorithms.

Fair MAA2C. It is a straight forward application of the policy gradient theorem. Each agent's actor is updated using the individual Actor's loss given by $\mathbb{E}_t [A_{i,t}^{F,\bar{\theta}} \log \pi_{\theta_i}(a_{i,t} | s_t)]$.

²Code available at <https://github.com/AkuBrains/altruistic-fair-MARL/>

Fair MAPPO. It is derived from MAPPO. We keep the core idea of MAPPO and we replace the classical advantage function with the fair advantage function in the PPO loss despite mathematical inaccuracies.

To assess the performance of our algorithms, we use two metrics. The first one evaluates the efficiency and is defined by the total number of apples consumed over an episode. The second one assesses the fairness among the agents and is defined by the Gini coefficient $Gini_t$ over the apples consumed, where

$$Gini_t(c_0, \dots, c_n) = \frac{\sum_{i \in N} \sum_{j \in N} |c_i - c_j|}{2N \sum_{i \in N} c_i}$$

with c_i being the number of apples consumed by agent i in an episode. The Gini coefficient navigates between 0.0 and 1.0. The lower it is, the fairer the distribution is, reaching 0.0 for an even distribution. On the other hand, the higher it is, the more unequal the distribution is.

We also study the fully cooperative setting in more depth and compare the Proportional Fairness (PF) objective and the Utilitarian Welfare (UW) objective.

5.3 Results

Figure 2 presents a comparison between the PF objective and the UW one in a fully cooperative settings ($\alpha = 1$). Surprisingly, the Prop. Fairness approach demonstrates superior overall efficiency. The total number of apples consumed by the group steadily increases, reaching approximately 120 by the end of the training. On the contrary, the Utilitarian driven method yields a significantly lower total apple consumption, leveled off near 40.

After breaking down the overall consumption, all seven agents consume a relatively similar number of apples in the PF setting. Although there are minor variations, the trend lines for all agents are clustered together. No single agent is left behind, and all seem to contribute and benefit from the fair task which is reflected in the low Gini coefficient. In stark contrast, only two agents (2 and 3) learn to harvest apples while the others are cleaning the river or doing nothing in the UW setting. Thus, this outcomes in a very unequal situation characterized by a very high Gini index near 0.8.

In the Utilitarian settings, each agent benefits from others' reward signal independently of the action the agent takes. Therefore, the reward gained by the agent and its actions are not necessarily correlated. Agents are trapped into sub-optimal policies creating

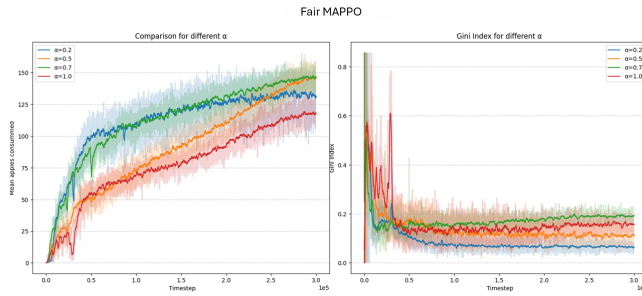


Figure 3: Performance of varying α in the CleanUp environment trained with Fair-MAPPO.

specialized role for each agent. When agents share the same parameters, they basically learn a single policy assuming natural cooperation and homogeneous behaviors. In the PF setting, the natural fairness brought by the PF objective encourages agents to explore further whenever an imbalance occurs. In the Fair Advantage function, the classical advantage function is normalized by the state-value function significantly reducing the weight of others’ reward signal in the PF objective.

Figure 3 shows the results for Fair MAPPO. All agent groups successfully learn to collect more apples over time, as indicated by the upward trend of all curves. Moreover, the Gini index of all experiments is quite low and all runs are clustered between 0.05 and 0.2. The group with $\alpha = 0.7$ achieves the best overall performance, consuming the most apples by the end of the training. Against all odds, the worst group is the fully cooperative one achieving significantly lower scores in both overall performance and Gini index.

Intuitively, in the fully cooperative setting, the same phenomenon as in the UW setting exists but attenuated by the equity willingness of the PF objective. In contrast, when $\alpha < 1$, agents are more inclined to follow their own objective and are prone harvesting apples autonomously. When α is too low ($\alpha = 0.2$), the overall harvesting performance drops. In that case, agents are merely selfish meaning they are less inclined to coordinate themselves to clean the river.

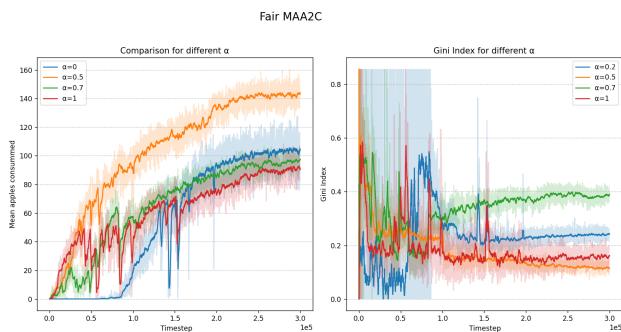


Figure 4: Performance of varying α in the CleanUp environment trained with Fair-MAA2C.

Figure 4 shows the results for Fair MAA2C. All agent groups successfully learn to collect more apples over time, as indicated by the upward trend of all curves. However, the training is very noisy as we can notice from the several steep drops of the average apples consumed during the training. It is well known that MAA2C training process is very unstable and our Fair MAA2C exhibits the same behavior making the results difficult to interpret. Nonetheless, the best utilitarian performance and Gini index is obtained for $\alpha = 0.5$ correlating Fair MAPPO results. In addition, Fair MAA2C seems to achieve lower mean apples consumed compared to its PPO counterpart except for $\alpha = 0.5$ probably due to instability during the learning process.

In short, selfish behavior can lead to a more equal outcome, but only because everyone is competing for the same individual rewards, potentially at the expense of the group’s overall success. This highlights the fundamental trade-off between individual selfishness and collective goals. We also demonstrate that the fairness promoted by the PF objective enables sustainable cooperation, whereas a utilitarian approach fails to do so, particularly in heterogeneous groups.

5.4 Discussions and Next steps

Our algorithms still need to be improved regarding theoretical proofs. In particular, there is no convergence proof to any kind of equilibrium if it exists. Moreover, replacing the traditional Advantage function with the Fair Advantage function is inconsistent with the theory as we didn’t prove the monotonic improvement theorem of Trust Region Policy Optimization in that case [24, 29].

On the experimental part, it would be interesting to test our algorithms in more environments [1, 11] or realistic real-world simulation [28] and create more diverse scenarii for CleanUp. With an apple regeneration rate of 0.05, the harvesting pressure exerted by only 7 agents is too low to create a significant imbalance between supply and demand. An pertinent scenario would be to increase the number of agents and decrease the regeneration rate to create a situation of resource scarcity, where demand is greater than supply.

Another way to improve our experiment in the Utilitarian setting would be to use the most recent state-of-the-art algorithms like HAPPO or HAA2C [18].

6 CONCLUSION

This work introduces a novel framework that integrates Proportional Fairness with altruism to foster more equitable cooperation in multi-agent systems. We proposed a fair altruistic utility and extended it to sequential settings by defining a Fair Altruistic Markov Game, from which we derived novel fair policy gradient algorithms.

Our experiments in the ‘CleanUp’ social dilemma demonstrate that this approach not only achieves significantly fairer outcomes, as measured by the Gini index, but also leads to superior overall efficiency compared to the traditional utilitarian welfare objective with heterogeneous agents. By effectively balancing individual incentives with collective well-being, our method prevents agents from adopting highly specialized and inequitable roles, thus paving the way for more stable and sustainable cooperation.

REFERENCES

- [1] John P. Agapiou, Alexander Sasha Vezhnevets, Edgar A. Duéñez-Guzmán, Jayd Matyas, Yiran Mao, Peter Sunehag, Raphael Köster, Udari Madhushani, Kavya Kopparapu, Ramona Comanescu, D. J. Strouse, Michael B. Johanson, Sukhdeep Singh, Julia Haas, Igor Mordatch, Dean Mobbs, and Joel Z. Leibo. 2023. Melting Pot 2.0. <https://doi.org/10.48550/arXiv.2211.13746> arXiv:2211.13746 [cs].
- [2] Farinaz Alamiyan-Harandi and Pouria Ramazi. 2023. Environmental-Impact Based Multi-Agent Reinforcement Learning. <https://doi.org/10.48550/arXiv.2311.04240> arXiv:2311.04240 [cs].
- [3] Stefano V. Albrecht, Filippos Christianos, and Lukas Schäfer. 2024. *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*. MIT Press. <https://www.marl-book.com>
- [4] Siddhartha Banerjee, Vasilis Gkatzelis, Safwan Hossain, Billy Jin, Evi Micha, and Nisarg Shah. 2022. Proportionally Fair Online Allocation of Public Goods with Predictions. <https://doi.org/10.48550/arXiv.2209.15305> arXiv:2209.15305 [cs].
- [5] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation* 7, 3 (Aug. 2019), 1–32. <https://doi.org/10.1145/3355902>
- [6] Po-An Chen, Bart De Keijzer, David Kempe, and Guido Schäfer. 2011. The Robust Price of Anarchy of Altruistic Games. In *Internet and Network Economics*, Ning Chen, Edith Elkind, and Elias Koutsoupias (Eds.). Vol. 7090. Springer Berlin Heidelberg, Berlin, Heidelberg, 383–390. https://doi.org/10.1007/978-3-642-25510-6_33 Series Title: Lecture Notes in Computer Science.
- [7] Robyn M. Dawes and David M. Messick. 2000. Social Dilemmas. *International Journal of Psychology* 35, 2 (2000), 111–116. <https://doi.org/10.1080/002075900399402> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1080/002075900399402>
- [8] Ishan Durugkar, Elad Liebman, and Peter Stone. 2020. Balancing Individual Preferences and Shared Objectives in Multiagent Reinforcement Learning. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, Yokohama, Japan, 2505–2511. <https://doi.org/10.24963/ijcai.2020/347>
- [9] Soroush Ebadian, Rupert Freeman, and Nisarg Shah. 2024. Harm Ratio: A Novel and Versatile Fairness Criterion. In *Proceedings of the 4th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. ACM, San Luis Potosi Mexico, 1–14. <https://doi.org/10.1145/3689904.3694701>
- [10] Promise Osaine Ekpo, Brian La, Thomas Wiener, Saesha Agarwal, Arshia Agrawal, Gonzalo Gonzalez-Pumariega, Lekan P. Molu, and Angélique Taylor. 2025. Skill-Aligned Fairness in Multi-Agent Learning for Collaboration in Healthcare. <https://doi.org/10.48550/arXiv.2508.18708> arXiv:2508.18708 [cs].
- [11] Zihao Guo, Shuqing Shi, Richard Willis, Tristan Tomilin, Joel Z. Leibo, and Yali Du. 2025. SocialJax: An Evaluation Suite for Multi-agent Reinforcement Learning in Sequential Social Dilemmas. <https://doi.org/10.48550/arXiv.2503.14576> arXiv:2503.14576 [cs].
- [12] Garrett Hardin. 1968. The tragedy of the commons: the population problem has no technical solution; it requires a fundamental extension in morality. *science* 162, 3859 (1968), 1243–1248.
- [13] Edward Hughes, Joel Z. Leibo, Matthew G. Phillips, Karl Tuyls, Edgar A. Duéñez-Guzmán, Antonio García Castañeda, Iain Dunning, Tina Zhu, Kevin R. McKee, Raphael Koster, Heather Roff, and Thore Graepel. 2018. Inequity aversion improves cooperation in intertemporal social dilemmas. arXiv:1803.08884 [cs.NE] <https://arxiv.org/abs/1803.08884>
- [14] Jiechuan Jiang and Zongqing Lu. 2019. Learning Fairness in Multi-Agent Systems. In *Advances in Neural Information Processing Systems*, Vol. 32. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2019/hash/10493aa88605cad5ab4752b04a63d172-Abstract.html
- [15] Peizhong Ju, Arnob Ghosh, and Ness B. Shroff. 2023. Achieving Fairness in Multi-Agent Markov Decision Processes Using Reinforcement Learning. <https://doi.org/10.48550/arXiv.2306.00324> arXiv:2306.00324 [cs].
- [16] Mamoru Kaneko and Kenjiro Nakamura. 1979. The Nash Social Welfare Function. *Econometrica* 47, 2 (March 1979), 423. <https://doi.org/10.2307/1914191>
- [17] FP Kelly, AK Maulloo, and DKH Tan. [n.d.]. Rate control for communication networks: shadow prices, proportional fairness and stability. ([n.d.]).
- [18] Jakub Grudzien Kuba, Xidong Feng, Shiyao Ding, Hao Dong, Jun Wang, and Yaodong Yang. 2022. Heterogeneous-Agent Mirror Learning: A Continuum of Solutions to Cooperative MARL. <https://doi.org/10.48550/arXiv.2208.01682> arXiv:2208.01682 [cs].
- [19] Weifan Long, Wen Wen, Peng Zhai, and Lihua Zhang. 2024. Role Play: Learning Adaptive Role-Specific Strategies in Multi-Agent Interactions. <https://doi.org/10.48550/arXiv.2411.01166> arXiv:2411.01166 [cs].
- [20] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. 2020. Multi-Agent Actor-Critic for Mixed Cooperative-Competitive Environments. <https://doi.org/10.48550/arXiv.1706.02275> arXiv:1706.02275 [cs].
- [21] Debmalya Mandal and Jiarui Gan. 2023. Socially Fair Reinforcement Learning. <https://doi.org/10.48550/arXiv.2208.12584> arXiv:2208.12584 [cs].
- [22] Hervé Moulin. 2003. *Fair Division and Collective Welfare*. The MIT Press. <https://doi.org/10.7551/mitpress/2954.001.0001>
- [23] Stefan Roesch, Stefanos Leonardos, and Yali Du. [n.d.]. Selfishness Level Induces Cooperation in Sequential Social Dilemmas. ([n.d.]).
- [24] John Schulman, Sergey Levine, Philipp Moritz, Michael I. Jordan, and Pieter Abbeel. 2017. Trust Region Policy Optimization. arXiv:1502.05477 [cs.LG] <https://arxiv.org/abs/1502.05477>
- [25] Umer Siddique, Peilang Li, and Yongcan Cao. 2024. Towards Fair and Equitable Policy Learning in Cooperative Multi-Agent Reinforcement Learning. (2024).
- [26] Richard S. Sutton and Andrew Barto. 2014. *Reinforcement learning: an introduction* (nachdruck ed.). The MIT Press, Cambridge, Massachusetts.
- [27] Chao Yu, Akash Velu, Eugene Vinitzky, Jiaxuan Gao, Yu Wang, Alexandre Bayen, and Yi Wu. 2022. The Surprising Effectiveness of PPO in Cooperative, Multi-Agent Games. <https://doi.org/10.48550/arXiv.2103.01955> arXiv:2103.01955 [cs].
- [28] Tianyu Zhang, Andrew Williams, Soham Phade, Sunil Srinivasa, Yang Zhang, Prateek Gupta, Yoshua Bengio, and Stephan Zheng. 2022. AI for Global Climate Cooperation: Modeling Global Climate Negotiations, Agreements, and Long-Term Cooperation in RICE-N. *SSRN Electronic Journal* (2022). <https://doi.org/10.2139/ssrn.4189735>
- [29] Yifan Zhong, Jakub Grudzien Kuba, Xidong Feng, Siyi Hu, Jiaming Ji, and Yaodong Yang. [n.d.]. Heterogeneous-Agent Reinforcement Learning. ([n.d.]).