

Learning in Multi-Objective Public Goods Games with Non-Linear Utilities

Nicole Orzan^{a,*}, Erman Acar^b, Davide Grossi^{a,b}, Patrick Mannion^c and Roxana Rădulescu^{d,e}

^aUniversity of Groningen

^bUniversity of Amsterdam

^cUniversity of Galway

^dUtrecht University

^eVrije Universiteit Brussels

ORCID (Nicole Orzan): <https://orcid.org/0000-0002-9204-0688>, ORCID (Erman Acar):

<https://orcid.org/0000-0001-7541-2999>, ORCID (Davide Grossi): <https://orcid.org/0000-0002-9709-030X>, ORCID

(Patrick Mannion): <https://orcid.org/0000-0002-7951-878X>, ORCID (Roxana Rădulescu):

<https://orcid.org/0000-0003-1446-5514>

Abstract. Addressing the question of how to achieve optimal decision-making under risk and uncertainty is crucial for enhancing the capabilities of artificial agents that collaborate with or support humans. In this work, we address this question in the context of Public Goods Games. We study learning in a novel multi-objective version of the Public Goods Game where agents have different risk preferences, by means of multi-objective reinforcement learning. We introduce a parametric non-linear utility function to model risk preferences at the level of individual agents, over the collective and individual reward components of the game. We study the interplay between such preference modelling and environmental uncertainty on the incentive alignment level in the game. We demonstrate how different combinations of individual preferences and environmental uncertainty sustain the emergence of cooperative patterns in non-cooperative environments (i.e., where competitive strategies are dominant), while others sustain competitive patterns in cooperative environments (i.e., where cooperative strategies are dominant).

1 Introduction

How can cooperation emerge and sustain itself in situations where agents do not necessarily have a direct motive for cooperation? This is one of the fundamental questions in various research areas, such as evolutionary biology [21, 35], political sciences [8, 9], cognitive sciences [43] and physics [11]. To answer this question, researchers developed and studied models of real-world scenarios involving tension between the collective and personal motives, called social dilemmas [13, 26]. The main characteristic of these social dilemmas is that players are better off defecting at the individual level, while, at the group level, the best outcome is mutual cooperation.

In this work, we focus on a specific class of social dilemmas known as Public Goods Games (PGG), extensively studied in literature [3, 47]. A PGG describes situations where cooperation by all agents is Pareto optimal, but because of the profitability of free-riding [5], rational agents fail to cooperate: defection by all agents

is a Nash equilibrium [34]. We refer to this kind of game as mixed-motives, since the incentives of the agents are partially misaligned.

In addition to incentive misalignment, other factors influencing the emergence of cooperation in many real-world scenarios include uncertainty and different individual attitudes towards risk [20, 27].

Uncertainty can have different sources: we refer to *environmental uncertainty* when actors are unsure about the amount of goods they can receive from the environment [4, 53], and to *social uncertainty* when it comes to ambiguity about the opponents' possible actions [10, 16]. Individual preferences express a personal inclination towards one choice over another. In the specific context of PGGs, we are interested in modelling two main types of individuals in the presence of uncertainty: those that are biased towards taking risks in the presence of uncertainty, also called *risk-seeking agents*, and others which are inclined not to take risks, also called *risk-averse agents*. We model these attitudes using a parametric nonlinear utility function of the reward received by individuals as the result of their investment in the collective good.

Since we are working with non-linear utility functions in the PGG, we need to decouple our perspective from the literature on the PGGs addressing non-linear public good productions. This branch focuses on settings where the public good results from a non-linear production process [41]. These are called non-linear public good games and allow one to model certain real-world situations (populations of bacteria, viruses, or cooperative hunting [12, 39]). In contrast, we shift our focus to an individual level and capture settings where potentially different attitudes towards risk can occur within a population.

To model risk attitudes, in our work, we explicitly decouple the collective versus the individual incentives experienced by the agents, and parameterize the collective incentive at the individual level. This choice allows us to model settings where individuals in a population can have different perceptions regarding these incentives. Furthermore, we take a multi-objective approach to the optimization of these two levels of rewards, drawing on multi-objective reinforcement learning (MORL) methods. This allows us to investigate learned behaviours that emerge from individually preferred trade-offs between the cooperative and competitive objectives.

* Corresponding Author. Email: n.orzan@rug.nl

Contributions. We investigate learning in PGGs where agents have different risk preferences, modelled as non-linearities over the game payoffs. We study the interplay between this mechanism and environmental uncertainty from a multi-objective perspective. More specifically, our contributions are as follows:

1. We propose a novel multi-objective multi-agent (MOMA) environment based on the extended Public Goods Game (EPGG) [38], called the Multi-Objective EPGG (MO-EPGG). This environment, next to facilitating training agents on games with different levels of incentive alignment, also allows one to explicitly model the trade-off between the individual and the cooperative components of PGGs. Moreover, it enables the decoupling of environmental and social uncertainties, allowing for the analysis of their impact, both when occurring concurrently or in isolation.
2. We propose a non-linear utility function that allows one to combine the collective and individual rewards, parameterized at the agent level. The selected shape of the utility function allows us to model risk-averse and risk-seeking agents, by operating a convex or concave transformation over the collective game reward, respectively. Moreover, it allows us to model a population of agents with varying attitudes towards risk.
3. We perform an analysis of the MO-EPGG under different multi-objective optimisation criteria. In particular, we look at the joint impact of different risk preferences and incentive alignment levels on the best responses, the Nash equilibria and the price of anarchy.
4. We perform experiments on a population of independent multi-objective reinforcement learning agents trained on the MO-EPGG. We show that risk-averse utility functions strongly diminish cooperation in cases with and without uncertainty. In the presence of environmental uncertainty, risk-seeking utilities improve cooperation in environments where defection is the dominant strategy.

2 Related Work

The related literature can be grouped under two main categories: non-linear utilities in public good games and multi-objective reinforcement learning. In this section, we describe them respectively.

2.1 Non-Linear Utilities in Public Goods Games

Although the PGG with linear utility functions is the most known and commonly used, various models of non-linear PGGs have been proposed in the literature as well. In the *threshold public goods game*, the resulting public good is given by a step function of the number of cooperators: the resource is created only if a minimum fraction of actors participate in the production of the public good [14]. When the minimum number of participants is 1, this is called the *Volunteer's Dilemma* [6, 17]. A *sigmoid public goods* function closely models many biological systems where the output production is small for low input levels and bigger for intermediate inputs, decreasing again for even bigger ones [7, 12]. In other paradigms, public good production is modelled by applying a concave (convex) function over agents' contributions, where the produced good is lesser (greater) than the good provided by a linear function of the contributions.

Several papers focused on analyzing non-linear public good games, by different means. In [33] authors employ non-linear PGG with different incentive structures to analyze behavioural subtyping, i.e., if cooperative behaviour in one task can predict cooperative behaviour in another. In [55], evolutionary dynamics techniques are employed to study the role of different non-linear production functions on the evolution of cooperation in finite populations, while in

[41], the evolutionary dynamics of two different populations collaborating for the production of a non-linear public good is investigated. In [15] authors explore the effects of different non-linear PGGs on the evolution of cooperation using Darwinian dynamics.

In the aforementioned literature, non-linearities in PGGs are typically functions that influence the production of the public good. In our work, however, we take a different perspective by introducing non-linearities at the level of the individual utilities extracted from rewards. More specifically, our goal is to model individuals' attitudes towards risk. In doing so, we follow decision theory which seeks to understand human decision processes and derive optimal decision-making strategies [51, 29, 24], therefore the study of risk and uncertainty has been a central focus. Some studies have shown that people make decisions based on some subjective function of the investment they made [48, 18]. For instance, an individual's risk attitude is often described as a function of the investment made (x) by means of a utility function shaped as $u(x) = x^\beta$. Here, the parameter β governs the risk preference of the individual: if $0 < \beta < 1$, the function is concave, signifying risk-aversion; if $\beta > 1$, the function is convex, indicating a risk-seeking attitude [24]. In our work, we draw on this idea to formulate a utility function that allows us to model individual preferences for actors participating in the PGG.

2.2 Multi-Objective Reinforcement Learning

In the field of reinforcement learning, the main focus is often to solve single-objective problems, by determining the agent's best policy to reach a specific goal. However, real-world challenges are of a multi-objective nature most of the time [45]. Autonomous agents, whether human or artificial, need to optimize for multiple goals simultaneously, or find a trade-off between them. This is the central concept of multi-objective reinforcement learning (MORL) [23, 45], a field that has developed rapidly in recent years [2, 54]. In MORL, the core idea is to receive vector rewards from the environment instead of scalar rewards. Under the utility-based perspective [23], rewards can be combined by means of a scalarization function to determine the final optimisation goal. Often, a linear scalarization function is employed, which allows the employment of single-objective RL methods. Alternatively, other choices include monotonically increasing non-linear scalarization functions [1, 44]. These are of particular interest for our work since non-linear functions are often used to model utilities under uncertainty and risk, especially in the economics literature, which aims at modelling human behaviour [33, 50].

Another part of this field of research focuses on fairness, i.e., how to optimize the trade-off among the objectives of different individuals under particular fairness constraints [49, 22, 19]. For example, in [49], authors employ deep RL techniques to learn a policy that treats users equitably. We build on the framework developed in [49], but rather than focusing on the fair treatment of a set of users, we investigate the effect of uncertainty and individuals' attitudes towards risk. To this end, we extend their approach to work with a different scalarisation function customized for our scenario, which allows us to model individual preferences, and train independent reinforcement learning agents in a multi-objective setting. We thus adopt a multi-objective multi-agent reinforcement learning (MOMARL) [42] perspective, which extends MORL to multi-agent scenarios.

3 Preliminaries

This section presents the formal definitions and the background knowledge. These include the Extended Public Goods Game, the

multi-objective stochastic games and the multi-objective optimization criteria.

3.1 The Extended Public Goods Game

The Extended Public Goods Games [36] is a tuple $\langle N, \mathbf{c}, A, f, \mathbf{u} \rangle$, where N is the set of players whose size is denoted as $|N| = n \in \mathbb{N}$. Every player i is endowed with some amount of wealth (or coins) $c_i \in \mathbb{R}_{\geq 0}$, and $\mathbf{c} = (c_1, \dots, c_n)$ denotes the tuple containing all agents' coins. Each agent can decide whether to invest in the public good (cooperate) or keep the endowment for themselves (defect); therefore, the set A of allowed actions consists of cooperate (C) and defect (D) i.e., $A = \{C, D\}$. The vector $\mathbf{a} = (a_1, \dots, a_n) \in A^n$ represents the action profile of the agents. The quantity f is called *multiplication factor*, and specifies the scalar by which the total investment is multiplied in order to produce the public good. The resulting quantity is then evenly distributed among all agents. The difference with respect to the original PGG lies in the interval of allowed values for f . While in the PGG $f \in (1, n)$, in the EPGG we take $f \in \mathbb{R}_{\geq 0}$. The reward function for each agent i is defined as $r_i : A^n \times \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}$, with:

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

where a_j is the j -th entry of the action profile \mathbf{a} and $I(a_j)$ is the indicator function, equal to 1 if the action of the agent j is cooperative, and 0 otherwise, and c_j denotes the j -th entry of \mathbf{c} . For the sake of simplicity, in the following, we assume all endowments to be equal, namely $c_i = c, \forall i \in N$.

Depending on the value of f , the EPGG can model three types of scenarios. When $1 < f < n$, like in the classic PGG, we model mixed-motives scenarios, in which all agents playing defect is a dominant strategy equilibrium. Yet, this profile is Pareto dominated by the profile in which all agents cooperate. When $0 \leq f \leq 1$, playing defect is a Pareto optimal dominant strategy (and therefore Nash) equilibrium. In addition, the EPGG can also model fully cooperative scenarios (i.e., when $f \geq n$) in which the cooperation profile is a Pareto optimal dominant strategy (and therefore Nash) equilibrium.

3.2 Multi-Objective Stochastic Games

We model the multi-objective multi-agent interactions using the *multi-objective stochastic game* (MOSG) framework, defined as the tuple $\mathcal{M} = (S, \mathcal{A}, T, \gamma, \mathcal{R})$, with $n \geq 2$ agents and $d \geq 2$ objectives, where:

- S is the state space,
- $\mathcal{A} = A_1 \times \dots \times A_n$ is the set of joint actions, with A_i being the action set of agent i ,
- $T : S \times \mathcal{A} \times S \rightarrow [0, 1]$ is the probabilistic transition function,
- γ is the discount factor,
- $\mathcal{R} = \mathbf{R}_1 \times \dots \times \mathbf{R}_n$ are the reward functions, where $\mathbf{R}_i : S \times \mathcal{A} \times S \rightarrow \mathbb{R}^d$ is the vectorial reward function of agent i for each of the d objectives.¹

We take a utility-based perspective [45] for multi-objective decision making, assuming that each agent i has a utility function $u_i : \mathbb{R}^d \rightarrow \mathbb{R}$ that maps the received reward vector to a scalar value, determining the desired trade-off between the objectives.

¹ We note that in this article the terms *reward* and *payoff* are synonyms. For the sake of clarity and consistency, we stick to the former term which aligns with the reinforcement learning terminology.

3.3 Optimization Criteria

In reinforcement learning, the goal of an agent is to find a policy π that maximizes the expected scalar return $V^\pi = \mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t r_t]$. In MORL, depending on how agents derive their utility, there are two optimisation criteria one can employ in the scalarisation process when maximising the expected discounted long-term reward vector:

- The Scalarised Expected Return (SER) criterion:

$$V_u^\pi = u \left(\mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t \mathbf{r}_t \right] \right), \quad (2)$$

where $\pi : S \times A \rightarrow [0, 1]$ is the agent's policy, and $\mathbf{r}_t = \mathbf{R}(s_t, a_t, s_{t+1})$ is the vectorial reward at timestep t .

- The Expected Scalarised Return (ESR) criterion:

$$V_u^\pi = \mathbb{E}_\pi \left[u \left(\sum_{t=0}^{\infty} \gamma^t \mathbf{r}_t \right) \right] \quad (3)$$

Which one of these criteria to choose depends on the problem at hand [42]. If we care about the goodness of a single policy execution, ESR is the correct criterion. If instead, we are interested in the quality of average policy executions, we should use SER. In this work, we opt for the SER criterion in the learning, modelling agents that are interested in optimising their behaviour in repeated interaction settings.

4 Multi-Objective EPGG

We formulate a multi-objective version of the EPGG, called Multi-Objective Extended Public Goods Game (MO-EPGG), by employing the framework of multi-objective stochastic games, outlined in Section 3.2. In our framework, the state space consists of the value of the multiplication factor f of the game currently being played, and the action space coincides with that of the single-objective EPGG (Section 3.1). We notice that the transition function for this framework is simply a random sampling from the set of possible multiplication factors at the beginning of each episode and deterministically returns that same f value at all the subsequent steps of the episode.

To complete our multi-objective formulation of the EPGG, we need to vectorize the scalar reward signal obtained by agents in the EPGG. This process is called *multi-objectivization* of single-objective problems [25, 30]. By observing the form of the reward function in Equation 1, we can easily distinguish between the part that defines the collective (r^C) and the individual payoff (r^I):

$$r_i^C(\mathbf{a}, f, \mathbf{c}) = \frac{1}{n} \sum_{j=1}^n c_j I(a_j) \cdot f \quad (4)$$

$$r_i^I(\mathbf{a}, \mathbf{c}) = c_i(1 - I(a_i)). \quad (5)$$

Then, in the proposed MO-EPGG, the vectorial reward received by agent i , given action profile \mathbf{a} , current multiplication factor f , and a tuple of endowments \mathbf{c} , is as follows:

$$\mathbf{r}_i(\mathbf{a}, f, \mathbf{c}) = (r_i^C(\mathbf{a}, f, \mathbf{c}), r_i^I(\mathbf{a}, \mathbf{c})). \quad (6)$$

This completes our description of the MO-EPGG as a MOSG with $d = 2$ objectives. In Figure 1 we display an example of the vectorial rewards received by $N = 2$ agents playing the MO-EPGG for three different values of the multiplication factor f .

To define the agents' utility functions, we follow a similar approach to the incentive structure proposed by Mullett et al. [33], but

note that we employ the non-linear function at an individual level, to model agents' risk attitudes as preferences over the received vectorial payoff. In particular, in our model, the gain in utility obtained from the collective reward behaves non-linearly by means of an exponential function where β_i serves as an exponent. Therefore, we define the following non-linear utility function that specifies the final scalarised utility for the MO-EPGG agents:

$$u_i(\mathbf{g}_i) = (g_i^C)^{\beta_i} + g_i^I, \quad (7)$$

with β_i being a hyperparameter. In our setting, g^C and g^I represent expected returns (or expected discounted sums of rewards), i.e., $g = \sum_t \gamma^t r_t$. Note that we are employing expected returns rather than rewards since we are working under the SER criterion. In this equation, the parameter β_i governs the risk-seeking/averse behaviour of agent i towards the collective expected return g_i^C , namely a value $\beta = 1$ returns a linear utility function, while $\beta < 1$ generates a concave function (with $\beta > 0$) which models a risk-avoiding agent, while $\beta > 1$ generates a convex function which models a risk-seeking agent [33].

In Equation 7, the exponent is only applied over the collective reward component. This choice is motivated by our conceptualization of the collective reward as the result of a risky investment. The result depends on the value of the multiplication factor f , which might not be known with certainty; and the actions of the other players.

$f = 0.5$		Player 1	
		C	D
Player 0	C	[2, 0], [2, 0]	[1, 0], [1, 4]
	D	[1, 4], [1, 0]	[0, 4], [0, 4]

$f = 1.5$		Player 1	
		C	D
Player 0	C	[6, 0], [6, 0]	[3, 0], [3, 4]
	D	[3, 4], [3, 0]	[0, 4], [0, 4]

$f = 2.5$		Player 1	
		C	D
Player 0	C	[10, 0], [10, 0]	[5, 0], [5, 4]
	D	[5, 4], [5, 0]	[0, 4], [0, 4]

Figure 1: Multi-objective payoff matrices for a 2-players MO-EPGG with 4 coins for each player, and multiplication factors 0.5, 1.5 and 2.5, when taking the cooperative (C) or defective (D) actions.

4.1 Game Analysis

We analyse the MO-EPGG with the proposed utility function. We first analyse the MO-EPGG under ESR and SER, to examine the dynamics of best responses for different values of the game and utility function parameters. Second, we investigate the impact of different values of β and f on the set of Nash equilibria under SER.

ESR. From Equation 7, we can observe that the preference between the cooperative or defective behaviour in the MO-EPGG depends on the relationship between three values, namely, f , c and β . In particular, assuming a uniform value of β among the whole population of agents ($\beta_i = \beta$ for all $i \in N$), the collective cooperative action ($\mathbf{a}_C = (C, \dots, C)$) is preferred over the collective defective action ($\mathbf{a}_D = (D, \dots, D)$) by all the agents when

$r^C(\mathbf{a}_C, f, \mathbf{c})^\beta > r^I(\mathbf{a}_D, f, \mathbf{c})$, which is the case when $(cf)^\beta > c$. This relationship between the variables induces a shared preference over collective cooperative behaviour in otherwise defective scenarios (the cases when $f < 1$). In general, collective cooperation is preferred over collective defection whenever either of the following conditions holds:

$$\beta < \frac{\log(c)}{\log(cf)} \quad \text{if } 0 < cf < 1 \quad (8)$$

$$\beta > \frac{\log(c)}{\log(cf)} \quad \text{if } cf > 1. \quad (9)$$

In the same way, collective defection is preferred over collective cooperation whenever $(cf)^\beta < c$.

SER. From this point onwards, we focus on analysing a 2-player MO-EPGG, under the SER criterion. We determine the minimum cooperation level of the opponent for which the player's best response is to cooperate. We compute this value for different values of f and β . The results are displayed in Table 1. We note that the minimum cooperation level of the opponent should be strictly greater than the value presented in the table, for the best response to be cooperation.

We can observe that, for the game with competitive incentive alignment ($f = 0.5$), the best response for each player is to defect every time the value of $\beta \leq 2$ (i.e., the opponent's probability to cooperate cannot be > 1 , hence the condition is unattainable). If $\beta > 2$ the best response is cooperation whenever the strategy of the opponent is to cooperate with a probability bigger than the value presented in the table. For example, for $\beta = 3$, the best response is cooperation whenever the strategy of the opponent is to cooperate with a probability bigger than 0.6. For $\beta = 4$, the threshold moves to 0.4, and to 0.3 for $\beta = 5$. For both the games with $f = 1.0$ and the game with a mixed-motive incentive alignment, $f = 1.5$, the best response is to defect every time $\beta \leq 2$, and to cooperate otherwise. For a game with cooperative incentive alignment, the best response is to cooperate every time $\beta \geq 1$.

Table 1: Minimum value for the cooperation of the opponent for which the best response strategy of the player is cooperation, computed for different values of f and β .

$f \backslash \beta$	0.5	1	2	3	4	5	6
0.5	1.	1.	1.	0.6	0.4	0.3	0.2
1.0	1.	1.	0.	0.	0.	0.	0.
1.5	1.	1.	0.	0.	0.	0.	0.
2.5	1.	0.	0.	0.	0.	0.	0.

NE under SER. To complement the above results, we perform an analysis of the Nash equilibria for the MO-EPGG played by two agents under SER. Similar to [31], we sweep through the space of possible joint strategies and identify the Nash equilibria² for a set of games. We select values of f in $\{0.5, 1.0, 1.5, 2.0, 2.5, 3.0\}$, to cover a spectrum of competitive, mixed-motive, and cooperative games. Additionally, we let the value of β span over the interval $[0., 3.]$, to capture risk-averse, risk-neutral and risk-seeking tendencies. The results are presented in Figure 2, where we display, for one agent, the value of the probability of cooperation in the Nash equilibria strategy as a function of β . We provide the plot for one agent only, since the corresponding plot for the other agent is identical.

² To check whether a joint strategy is an equilibrium under SER, we use the `iterated_best_response` and `verify_nash` methods of the Ramo library [46].




Figure 2: The probability of action C for Player 0 under the NE of a 2-player MO-EPGG, for varying values of f and β . The corresponding plot for Player 1 is identical. We note that in the case in which two mixed-strategy NE are present for the same value of β , the joint strategies are formed by the two different points present. The strategies of the agents are identical for pure-strategy NE.

From Figure 2 we can observe, for all the games, the coexistence of two Nash equilibria whenever $\beta < 1$. At each of these equilibria, the joint strategy of the two agents displays the following symmetry: if one agent defects, they both are better off when the other agent slightly moves away from full defection. For $f < 2$ and $\beta > 1$ we observe the existence of an interval of β values in which both mutual defection and cooperation coexist as Nash equilibria. Based on the results in Table 1, we notice that for a high enough value of β , even for $f = 0.5$, the mutual defection Nash equilibrium will eventually disappear. Whenever $\beta > 1$ and $f > 2$, as expected, the only Nash equilibrium is mutual cooperation. We present additional results on the SER in the Supplementary Material [37].

4.2 Price of Anarchy

To evaluate the goodness of the possible outcomes of the system against the Nash equilibria of each game in the MO-EPGG (defined by a specific f value), we adapt the metric known as Price of Anarchy (PoA) [28, 40]. The PoA is the ratio between the welfare of the system in its “best solution”, i.e. social optimum, and the welfare of the system at its worst NE. Thus, it expresses the potential degradation factor of the social optimum. We highlight that higher values of the PoA indicate the existence of possible worst outcomes for the system due to agents’ selfishness, in comparison to the social optimum.

On the other hand, a PoA with a value of 1 indicates an overlap between the social optimum and the worst-case selfish action. To define the social optimum in our setting, we employ the utilitarian welfare function, summing the outcomes of the utility functions of the players³: $W(\pi) = \sum_{i=0}^n u_i(\mathbf{V}_i^\pi)$, where \mathbf{V}_i^π is the expected vectorial return of player i , under the joint strategy π . Then, the PoA is defined as follows:

$$PoA = \frac{\max_s W(\pi)}{\min_{\pi \in \text{Nash}} W(\pi)}. \quad (10)$$




Figure 3: Price of Anarchy for varying values of f and β .

In Figure 3 we display the values of the PoA for a 2-player MO-EPGG, varying the values of f and β . From the figure, we can observe that for every f , the value of the PoA is 1 for $\beta < 0.5$. This is due to mutual defection being the only Nash equilibrium and best welfare point. We can also observe that every f displays a range of β for which the value of the PoA is greater than 1. This is due to the presence of a Nash equilibria in mutual defection when the welfare is instead maximized in mutual cooperation. We can also observe that the PoA value stabilizes to 1 for every game when β overcomes a certain threshold, which is dependent on the value of f . This is due to mutual cooperation being both the strategy that maximizes the welfare and the only Nash equilibrium of the game.

5 Experimental setup

MO-DQN We train independent RL agents by adapting the multi-objective version of the Deep Q-network (DQN) algorithm [32] described in [49], which allows us to optimize policies under the SER criterion⁴ Siddique et al. [49] train a DQN to predict a Q-function for every objective. Therefore, the dimension of the output is $|A| \times d$. We adjust their approach to work with our scalarization function (Equation 7). The loss function for MO-DQN can be expressed as follows:

$$L(\theta) = \mathbb{E}_{s,a,s',r \sim D} \left[\left(r + \gamma \hat{Q}_{\theta'}(s', a^*) - \hat{Q}_{\theta}(s, a) \right)^2 \right], \quad (11)$$

where θ and θ' represent the DQN weights at two different timesteps of the training. D represents the buffer of stored transitions, and r is the vector reward. We find the best action a^* by applying the SER optimization criterion, namely, by applying our custom scalarization function u to update the MO-DQN function:⁵

$$a^* = \operatorname{argmax}_{a \in A} u \left(\mathbb{E} [r + \gamma \hat{Q}_{\theta'}(s', a')] \right). \quad (12)$$

³ In the context of multi-objective games under SER, we cannot apply the welfare function directly on the payoffs of the matrix game.
⁴ We stress that the need for function approximation is due to the continuous input values provided by f .
⁵ Similar to [49], we compute the expectation of the scalarisation, which is a lower bound for the SER criterion.

5.1 Experiments

The experiments are run over a pool of $N = 20$ agents. At each iteration t of the learning process, a multiplication factor f_t is sampled uniformly from the interval $[f_{\min}, f_{\max}]$, where f_{\min} and f_{\max} are chosen such as to include cooperative, competitive, and mixed-motive games. Afterwards, a subset with $M = 4$ active agents is randomly sampled from the pool of N agents, to participate in the game for 10 consecutive rounds. After these interactions, the MO-DQN networks are updated. Given $M = 4$, we picked $f_{\min} = 0.5$ and $f_{\max} = 6.5$, to enable agents to engage in competitive, mixed-motive and cooperative games. This enables sampling from a set that contains competitive ($f < 1$), mixed-motive ($1 < f < M$), and cooperative games ($f > M$). Each agent receives as observation the current value of the multiplication factor—which can be observed with uncertainty—together with the previous actions taken by each opponent at the previous time step: $\mathbf{o}_t^i = (f_{obs}^i, \mathbf{a}_{t-1}^{-i})$, where $\mathbf{a}^{-i} = (a^j)_{j \in M, j \neq i}$. Therefore, each agent learns a policy $\pi_i : \mathcal{O}_i \times A_i \rightarrow [0, 1]$, where \mathcal{O}_i is the set of all possible observations of agent i .

We model uncertainty over the observation of the multiplication factor as Gaussian noise over the value of f , received from the environment: $f_{obs}^i = f + \mathcal{N}(0, \sigma_i^2)$, where σ_i is the uncertainty experienced by agent i . To ensure that the value of the observed multiplication factor coheres with the set of allowed values of f in the MO-EPGG, we round up every negative sampled value to 0.

All the experiments are run for 20000 epochs, and results are averaged over 20 runs for every condition. The *RMSprop* learning rate is set to $\lambda = 0.001$, and $\gamma = 0.99$. The action selection mechanism is ϵ -greedy, with $\epsilon = 0.01$. The values of the weights are $w^C = w^I = 1$ for all the agents. All the plots show the values of the average cooperation of the active agents at every evaluation step of the learning process. The DQN networks are composed of 2 hidden layers, and ReLU nonlinearities are employed between layers.⁶

6 Results

We group our empirical results in two categories, namely *homogeneous preferences*, where the value of β is identical for every agent, and *heterogeneous preferences*, where each agent i is characterised by a different β_i value. Both categories include experiments with and without uncertainty on the observation of the multiplication factor f .

6.1 Learning with homogeneous preferences

We first explore the impact of different values of β on the scenarios with and without uncertainty on the observations. We performed experiments for different values of β , that define a linear ($\beta = 1$), a convex ($\beta > 1$) and a concave ($\beta < 1$) utility function. In each of these experiments, β values are identical for every agent. The results for these experiments are depicted in Figure 4.

The experiments with $\beta = 1$ represent the baseline where agents are playing the linear version of the MO-EPGG. Therefore, in the games without uncertainty, we observe as expected convergence to cooperation whenever $f > M$, convergence to defection whenever $f < 1$, and a certain percentage of cooperation when $1 < f < M$, depending on whether the value of f is closer to a cooperative or a competitive one. When uncertainty is introduced, cooperation is increased in the competitive and mixed-motive scenarios. This results

from the concurrent learning on a set of games with different levels of incentive alignment, as previously observed in [38].

The experiments with $\beta = 0.5$ represent a system of risk-avoiding agents playing the MO-EPGG. This risk preference strongly pushes the system's behaviour toward competition across all games. The result stays consistent across the scenarios with and without uncertainty. This result is consistent with the analytical findings under SER outlined in Section 4.1: for all the games with different f values, (ϵ -)collective defection strategies are the only the Nash equilibria.

The experiments with $\beta \geq 2$ induce a system of risk-seeking agents playing in the MO-EPGG. Here, we observe that the cooperation of the system is drastically increased in all games with respect to the baseline $\beta = 1$. We note that for $\beta = 2$ and $f = 0.5$ agents do not converge to either cooperation or defection. This outcome aligns with the existence of both collective cooperation and collective defection as Nash equilibria.

6.2 Learning with heterogeneous preferences

Secondly, we investigate the impact of learning in the MO-EPGG when the agents' preferences β_i are heterogeneous, i.e., the value of β_i for every agent i is sampled from a normal distribution centred in 1, i.e. $\beta_i \sim \mathcal{N}(\mu_\beta, \sigma_\beta^2) \forall i \in N$, with $\mu_\beta = 1$. This allows us to get more values centred around risk neutrality, and few extreme risk-averse or risk-seeking tendencies. We performed experiments with different values of σ_β , i.e., 0.5, 2 and 3. The resulting system represents a population where every individual has a different risk preference and is centred on risk-neutrality ($\beta = 1$).

Figure 5 reports the results for the scenarios without (top row) and with (bottom row) uncertainty on the observations. We can observe that when the σ_β of the distribution is small (i.e. $\sigma_\beta = 0.5$) and no uncertainty is introduced, the competitive equilibria for $f \in \{0.5, 1.5\}$ is maintained, while the cooperative equilibria for $f \in \{3.5, 6.5\}$ is lost. We also notice that the higher the value of σ_β , the higher the average cooperation of the system in all the games. This result signals the importance of the magnitude of beta on the risk attitude: the higher β , the higher the risk appetite, which in our case translates to more cooperative behaviour.

When uncertainty is introduced, we observe that cooperation is increased in the competitive and mixed games with respect to the cases without uncertainty. This result is consistent with previous findings on the presence of uncertainty in non-cooperative environments [36, 38]. Interestingly, only the non-cooperative games are affected by the presence of uncertainty: in the cooperative games, the average cooperation of the system is equal to the one observed in the scenario without uncertainty.

6.3 Equilibria and Learning

We compare now our experimental results with the analysis from Section 4.1. We underline that, while the computation of equilibria is game-specific, the outcomes of the experiments result from concurrent learning on the set of environments modeled by the MO-EPGG.

Comparing the plots in Figure 4 and the analytical results of the Nash equilibria from Figure 2, we can observe that, in the case without uncertainty, in all the games for which $f \neq 0.5$, namely $f \in \{1.5, 3.5, 6.5\}$, the pool of agents learns the best response, which means that the system converges to the Nash equilibrium. Only for $\beta = 1$ the convergence is not perfect. In the case $f = 0.5$, the agents converge to defection when $\beta < 2$, while for $\beta \geq 2$ they cooperate with a certain probability, which is higher for higher β

⁶ All the parameters employed to perform the experiments are described in the Supplementary Material, in [37].

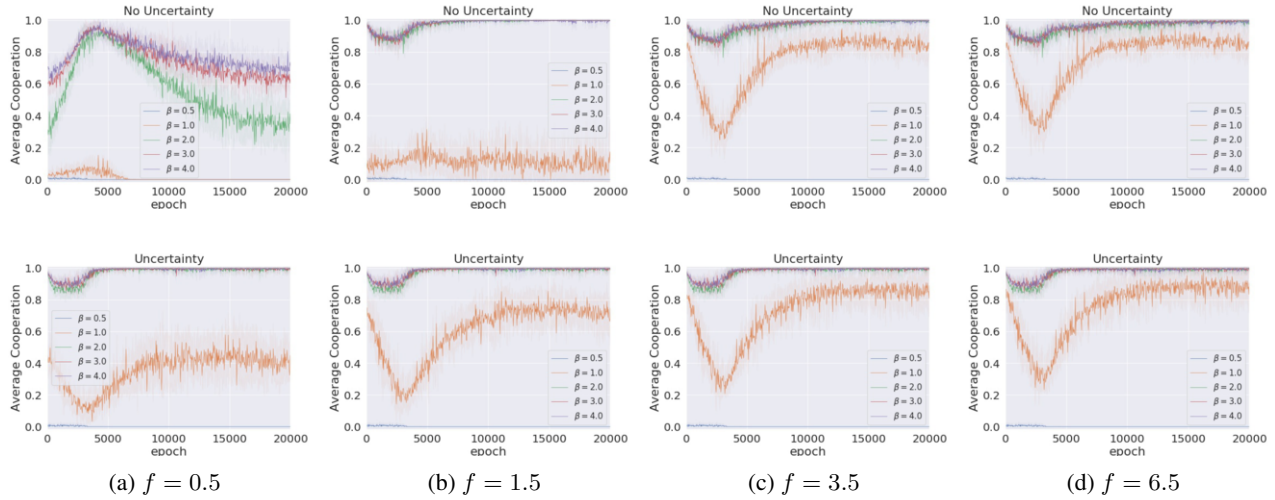


Figure 4: Average cooperation values for the active DQN agents trained across environments with different multiplication factors, without (top row) and with uncertainty (bottom row) on the observed multiplication factor, with $\sigma_i = 2, \forall i \in N$. The different values of β are identical for every agent $\beta_i = \beta, \forall i \in N$.

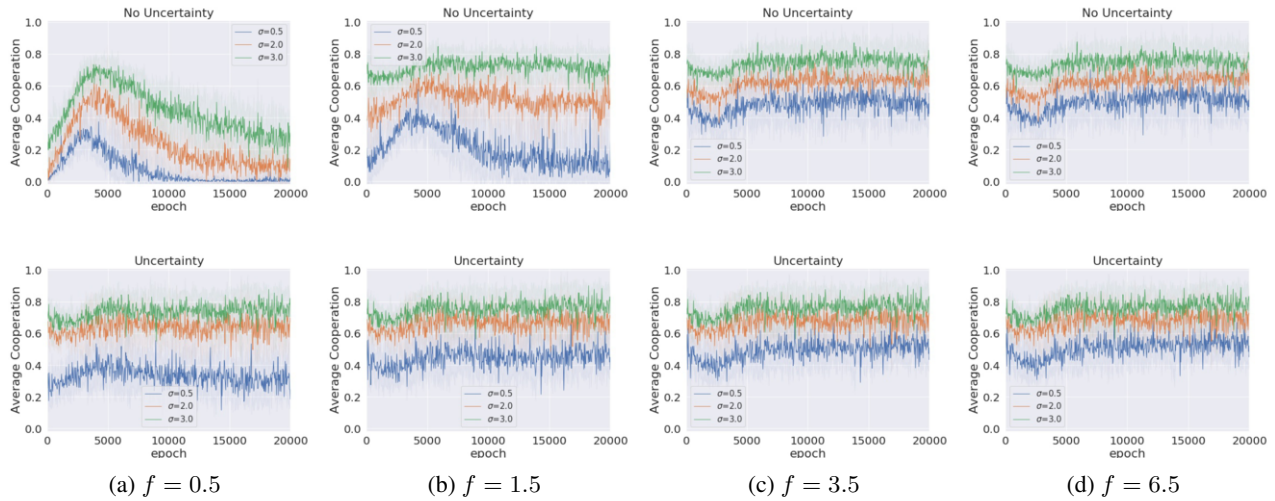


Figure 5: Average cooperation values for the active DQN agents trained across environments with different multiplication factors, without (top row) and with uncertainty (bottom row) on the observed multiplication factor, with $\sigma_i = 2 \forall i \in N$. The values of β are randomly sampled from a normal distribution $\beta_i \sim \mathcal{N}(\mu_\beta, \sigma_\beta^2) \forall i \in N$, with $\mu_\beta = 1$ and three different values of $\sigma_\beta = 0.5, 2, 3$.

values. This result is not surprising given the presence of more than one Nash equilibrium. When uncertainty is introduced, as previously mentioned, the average cooperation of the system increases.

7 Conclusions and Future Work

In this work, we introduced and analyzed a multi-objective variant of the Extended Public Goods Game, the MO-EPGG, which allowed us to model individual risk attitudes, by decoupling the collective and the individual payoffs. We investigated the role of misalignment of incentives, uncertainty and individual risk preferences on cooperation, utilising multi-objective reinforcement learning. In particular, we observed how risk-averse attitudes can increase defection in cooperative environments, and, inversely, risk-seeking ones, can increase cooperation in competitive and mixed-motive games, especially when uncertainty is introduced. Moreover, we observed how a population with heterogeneous risk attitudes, centred on risk neutral-ity, can fail to reach cooperation in cooperative settings.

Future work will explore additional MORL approaches, such as policy-gradient methods, in the MO-EPGG. We also plan to investigate the interplay between risk preferences, uncertainty, reputation mechanisms and social norms. Last but not least, we plan to explore adaptive risk attitudes [52] and their impact on the learning dynamics and cooperation levels in our setting.

Acknowledgements

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

References

[1] M. Agarwal, V. Aggarwal, and T. Lan. Multi-objective reinforcement learning with non-linear scalarization. In *Proceedings of the 21st Inter-*

- national Conference on Autonomous Agents and Multiagent Systems, pages 9–17, 2022.
- [2] L. N. Alegre, A. Bazzan, A. Nowe, and B. C da Silva. Multi-step generalized policy improvement by leveraging approximate models. *Advances in Neural Information Processing Systems*, 36, 2024.
 - [3] S. P. Anderson, J. K. Goeree, and C. A. Holt. A theoretical analysis of altruism and decision error in public goods games. *Journal of Public Economics*, 70(2):297–323, 1998.
 - [4] P. Andras, J. Lazarus, G. Roberts, and S. J. Lynden. Uncertainty and cooperation: Analytical results and a simulated agent society. *Journal of Artificial Societies and Social Simulation*, 2006.
 - [5] J. Andreoni. Why free ride?: Strategies and learning in public goods experiments. *Journal of public Economics*, 37(3):291–304, 1988.
 - [6] M. Archetti and I. Scheuring. Coexistence of cooperation and defection in public goods games. *Evolution*, 65(4):1140–1148, 2011.
 - [7] M. Archetti and I. Scheuring. Evolution of optimal hill coefficients in nonlinear public goods games. *Journal of Theoretical Biology*, 406: 73–82, 2016.
 - [8] R. Axelrod. The emergence of cooperation among egoists. *American political science review*, 75(2):306–318, 1981.
 - [9] R. Axelrod and W. D. Hamilton. The evolution of cooperation. *science*, 211(4489):1390–1396, 1981.
 - [10] J. Bendor. Uncertainty and the evolution of cooperation. *Journal of Conflict resolution*, 37(4):709–734, 1993.
 - [11] D. Challet and Y.-C. Zhang. Emergence of cooperation and organization in an evolutionary game. *Physica A: Statistical Mechanics and its Applications*, 246(3-4):407–418, 1997.
 - [12] J. S. Chuang, O. Rivoire, and S. Leibler. Cooperation and hamilton’s rule in a simple synthetic microbial system. *Molecular systems biology*, 6(1):398, 2010.
 - [13] R. M. Dawes. Social dilemmas. *Annual review of psychology*, 31(1): 169–193, 1980.
 - [14] K. De Jaegher. High thresholds encouraging the evolution of cooperation in threshold public-good games. *Scientific Reports*, 10(1):5863, 2020.
 - [15] K. Deng and T. Chu. Adaptive evolution of cooperation through darwinian dynamics in public goods games. *PLoS One*, 6(10):e25496, 2011.
 - [16] P. Deutchman, D. Amir, M. R. Jordan, and K. McAuliffe. Common knowledge promotes cooperation in the threshold public goods game by reducing uncertainty. *Evolution and Human Behavior*, 43(2):155–167, 2022.
 - [17] A. Diekmann. Volunteer’s dilemma. *Journal of conflict resolution*, 29 (4):605–610, 1985.
 - [18] J. Dow and S. R. da Costa Werlang. Uncertainty aversion, risk aversion, and the optimal choice of portfolio. *Econometrica: Journal of the Econometric Society*, pages 197–204, 1992.
 - [19] Z. Fan, N. Peng, M. Tian, and B. Fain. Welfare and fairness in multi-objective reinforcement learning. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, pages 1991–1999, 2023.
 - [20] E. Fehr and U. Fischbacher. Why social preferences matter—the impact of non-selfish motives on competition, cooperation and incentives. *The economic journal*, 112(478):C1–C33, 2002.
 - [21] S. A. Frank. Mutual policing and repression of competition in the evolution of cooperative groups. *Nature*, 377(6549):520–522, 1995.
 - [22] N. A. Grupen, B. Selman, and D. D. Lee. Cooperative multi-agent fairness and equivariant policies. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 9350–9359, 2022.
 - [23] C. F. Hayes, R. Rădulescu, E. Bargiacchi, J. Källström, M. Macfarlane, M. Reymond, T. Verstraeten, L. M. Zintgraf, R. Dazeley, F. Heintz, et al. A practical guide to multi-objective reinforcement learning and planning. *Autonomous Agents and Multi-Agent Systems*, 36(1):26, 2022.
 - [24] J. G. Johnson and J. R. Busemeyer. Decision making under risk and uncertainty. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(5): 736–749, 2010.
 - [25] J. D. Knowles, R. A. Watson, and D. W. Corne. Reducing local optima in single-objective problems by multi-objectivization. In *International conference on evolutionary multi-criterion optimization*, pages 269–283. Springer, 2001.
 - [26] P. Kollock. Social dilemmas: The anatomy of cooperation. *Annual review of sociology*, 24(1):183–214, 1998.
 - [27] S. Kopelman, J. M. Weber, and D. M. Messick. Factors influencing cooperation in commons dilemmas: A review of experimental psychological research. *The drama of the commons*, pages 113–156, 2002.
 - [28] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *Computer science review*, 3(2):65–69, 2009.
 - [29] H. Levy and M. Levy. Arrow-pratt risk aversion, risk premium and decision weights. *Journal of Risk and Uncertainty*, 25:265–290, 2002.
 - [30] X. Ma, Z. Huang, X. Li, Y. Qi, L. Wang, and Z. Zhu. Multiobjec-tivization of single-objective optimization in evolutionary computation: a survey. *IEEE Transactions on Cybernetics*, 2021.
 - [31] P. Mannion and R. Rădulescu. Comparing utility-based and pareto-based solution sets in multi-objective normal form games. In *Multi-Objective Decision Making Workshop 2023*, 2023.
 - [32] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Belle-mare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518 (7540):529–533, 2015.
 - [33] T. L. Mullett, R. L. McDonald, and G. D. Brown. Cooperation in public goods games predicts behavior in incentive-matched binary dilemmas: Evidence for stable prosociality. *Economic Inquiry*, 58(1):67–85, 2020.
 - [34] J. F. Nash Jr. Equilibrium points in n-person games. *Proceedings of the national academy of sciences*, 36(1):48–49, 1950.
 - [35] M. A. Nowak, A. Sasaki, C. Taylor, and D. Fudenberg. Emergence of cooperation and evolutionary stability in finite populations. *Nature*, 428 (6983):646–650, 2004.
 - [36] N. Orzan, E. Acar, D. Grossi, and R. Rădulescu. Emergent cooperation and deception in public good games. In *Adaptive and Learning Agents (ALA) Workshop at AAMAS*, 2023. URL <https://alaworkshop2023.github.io>.
 - [37] N. Orzan, E. Acar, D. Grossi, P. Mannion, and R. Rădulescu. Learning in multi-objective public goods games with non-linear utilities, 2024. URL <https://arxiv.org/abs/2408.00682>.
 - [38] N. Orzan, E. Acar, D. Grossi, and R. Rădulescu. Emergent cooperation under uncertain incentive alignment. In *Proceedings of the 2024 International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, 2024.
 - [39] C. Packer and L. Ruttan. The evolution of cooperative hunting. *The American Naturalist*, 132(2):159–198, 1988.
 - [40] C. Papadimitriou. Algorithms, games, and the internet. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 749–753, 2001.
 - [41] A. Patra, V. K. Dubey, and S. Chakraborty. Coexistence of coordination and anticoordination in nonlinear public goods game. *Journal of Physics: Complexity*, 3(4):045006, 2022.
 - [42] R. Rădulescu, P. Mannion, D. M. Roijers, and A. Nowé. Multi-objective multi-agent decision making: a utility-based analysis and survey. *Autonomous Agents and Multi-Agent Systems*, 34(1):10, 2020.
 - [43] D. G. Rand and M. A. Nowak. Human cooperation. *Trends in cognitive sciences*, 17(8):413–425, 2013.
 - [44] M. Reymond, C. F. Hayes, D. Steckelmacher, D. M. Roijers, and A. Nowé. Actor-critic multi-objective reinforcement learning for non-linear utility functions. *Autonomous Agents and Multi-Agent Systems*, 37(2):23, 2023.
 - [45] D. M. Roijers, P. Vamplew, S. Whiteson, and R. Dazeley. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48:67–113, 2013.
 - [46] W. Röpke. Ramo: Rational agents with multiple objectives. <https://github.com/wilrop/mo-game-theory>, 2022.
 - [47] F. C. Santos, M. D. Santos, and J. M. Pacheco. Social diversity promotes the emergence of cooperation in public goods games. *Nature*, 454(7201):213–216, 2008.
 - [48] L. J. Savage. *The foundations of statistics*. Courier Corporation, 1972.
 - [49] U. Siddique, P. Weng, and M. Zimmer. Learning fair policies in multi-objective (deep) reinforcement learning with average and discounted rewards. In *International Conference on Machine Learning*, pages 8905–8915. PMLR, 2020.
 - [50] A. Tversky and P. Wakker. Risk attitudes and decision weights. *Econometrica: Journal of the Econometric Society*, pages 1255–1280, 1995.
 - [51] J. Von Neumann and O. Morgenstern. Theory of games and economic behavior, 2nd rev. 1947.
 - [52] G. Wang, Q. Su, L. Wang, and J. B. Plotkin. The evolution of social behaviors and risk preferences in settings with uncertainty. *Proceedings of the National Academy of Sciences*, 121(30):e2406993121, 2024.
 - [53] A. Wit and H. Wilke. Public good provision under environmental and social uncertainty. *European journal of social psychology*, 28(2):249–256, 1998.
 - [54] R. Yang, X. Sun, and K. Narasimhan. A generalized algorithm for multi-objective reinforcement learning and policy adaptation. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*. Curran Associates Inc., Red Hook, NY, USA, 2019.
 - [55] Y. Zhang, F. Fu, T. Wu, G. Xie, and L. Wang. A tale of two contribution mechanisms for nonlinear public goods. *Scientific reports*, 3(1):2021, 2013.