Cooperation and Learning Dynamics under Risk Diversity and Financial Incentives

Ramona Merhej
Instituto Superior Tecnico, Portugal
ISIR, Sorbonne University
France
ramona.merhej@tecnico.ulisboa.pt

Fernando P. Santos
Informatics Institute
University of Amsterdam
Netherlands
f.p.santos@uva.nl

Francisco S. Melo
INESC-ID and Instituto Superior Tecnico, Universidade de Lisboa
Portugal
fmelo@inesc-id.pt

Mohamed Chetouani
ISIR, Sorbonne University
France
mohamed.chetouani@sorbonne-universite.fr

Francisco C. Santos
INESC-ID and Instituto Superior Tecnico, Universidade de Lisboa
Portugal
franciscocsantos@tecnico.ulisboa.pt

ABSTRACT

In this paper, we investigate the role of risk diversity in groups of agents learning to play collective risk dilemmas (CRDs). We show that risk diversity poses new challenges to cooperation that are not observed in homogeneous groups. While increasing average risk contributes, in general, for agents to cooperate with higher probability, increasing risk diversity significantly reduces a population's ability to achieve a collective target. Risk diversity leads to asymmetrical changes in agents policies — i.e. the increase in contributions from individuals at high risk is unable to compensate for the decrease in contributions from individuals at low risk — which reduces the total contributions in a population and overall social welfare. At the same time, risk diversity offers novel opportunities to design financial incentives, which, as we show, can improve cooperation, target achievement and global welfare beyond the levels obtained in the absence of diversity. Our results highlight the need to align risk perceptions among agents and implement diversity-based incentive policies in order to improve collectives' abilities to avoid future catastrophic events.

KEYWORDS

Social Dilemmas; Collective Risk Dilemmas; Cooperation; Risk Diversity; Financial Incentives; Heterogeneous Agents; Reinforcement Learning; Social Simulation

ACM Reference Format:
Ramona Merhej, Fernando P. Santos, Francisco S. Melo, Mohamed Chetouani, and Francisco C. Santos. 2022. Cooperation and Learning Dynamics under Risk Diversity and Financial Incentives. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022*, IFAAMAS, 9 pages.

1 INTRODUCTION

The World Economic Forum recently (January 2021) published its 16^{th} report on global risks [18]. Among the most concerning risks are climate change, biodiversity loss, extreme weather, as well as societal division and economic fragility. While some of these threats are on the verge of becoming unavoidable, individual and collective behaviors can still be adapted to avoid disastrous outcomes. This requires cooperation and coordination between a large number of agents, from citizens to policy-makers. Environmental, societal, economic, and other foreseeable collective risks require collective efforts — whether on a regional, national or international level — for successful resolutions of the problems and avoidance of crises.

While it is evident that large collective efforts are needed to avoid these disasters, people, institutions or countries remain reluctant to cooperate. On the one hand, no entity has the power of saving the system on its own. This is known as the problem of many hands (PMH) and occurs in interactions with a large number of players [55]. This challenge becomes even more prominent when actions are not directly harmful but only cause the risk of a harm [54]. On the other hand, cooperation in such contexts entails a social dilemma: the best individual outcome occurs if others contribute to the collective good and risks are avoided without one's intervention. This selfish reasoning, and the shifting of responsibility onto others, configures the so-called tragedy of the commons [22]. The tension within individuals/entities created by the urgent need of cooperation, the individually rational choice to defect, and the uncertainty about future outcomes and strategies of others, makes decision making non-trivial [3, 4, 25, 43]. The collective risk dilemma (CRD) is a simple game metaphor that tries to capture such challenges [11, 31, 40, 44, 56, 57].

In a CRD, agents decide how much of their wealth to contribute to a common pool with the goal of collectively achieving a target threshold that avoids future losses. If the target is not achieved, there is a risk that large losses are incurred by everyone.

The dynamics within groups playing CRDs have been tackled using different tools, from behavioral experiments [9, 11, 31, 51] to evolutionary game theory [40, 42, 44, 45, 57] and multi-agent reinforcement learning [12, 30]. However, previous works on CRDs considered a homogeneous risk factor [11, 31, 40, 45]. In reality, heterogeneous risk exposures and perceptions are ubiquitous. Recently, the COVID-19 crisis highlighted discrepancies in risk perception among countries [1, 20], as well as discrepancies in risk perception and exposure among individuals [7, 26, 34, 52]. Some studies have looked into other types of heterogeneities among agents and have reported significant changes in cooperation and target achievement

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online.

[23, 30, 57]. The findings on other heterogeneities motivated us to investigate the effect of introducing risk diversity in a population of agents facing collective risks. We examine how averaging out the risk value instead of considering risk diversity can alter the results we observe. To the best of our knowledge, this is the first work to explore risk diversities in the context of CRDs.

In multi-player games with common risks or common goods (i.e., where agents contribute to a common good instead of escaping a common risk), financial incentives (FI), in the form of punishments or rewards, are an effective mechanism to increase cooperation and target achievement [13, 21, 48–50]. After studying the effects of risk diversity, we explore how FI between peers at different risk exposure, hence with different interests for cooperation, can unify collective efforts and decrease the chances of failure. We investigate the benefits of FI on target achievement, overall cooperation and the secured welfare of agents. We contribute by showing how heterogeneity and non-symmetrical games can conceive new use cases for FI.

While the game tensions play a decisive role in the choices made by the agents, the final equilibrium of the system will also largely depend on the decision making process of agents. In fact, a game can be either studied from a static or a dynamic perspective based on how agents are assumed to make their decisions. A static perspective often assumes that agents are rational and fully know all possible strategy profiles and their respective outcomes which leads to convergence to Nash equilibria. Yet, experimental studies have shown that humans often do not select the rational choices [15, 19, 29, 47], and rather adapt their policies based on experience. Reinforcement learning (RL) suggests new tools to analyse decision making dynamics and, in fact, was shown to accurately model human behaviors in several social dilemmas [39]. Reinforcement learning has rapidly evolved in the past years, and several variations were developed specifically to promote cooperation in social dilemmas [16, 24]. However, the aforementioned algorithms require sharing large amounts of information and have therefore mostly been applied to 2-player games. The goal of this paper is to examine the cooperation challenges that emerge under RL dynamics in *large* player and *asymmetric* games with risk diversity. As such, we do not exploit algorithms designed to encourage cooperation in social dilemmas, but focus on independent RL where agents can only observe their own actions and rewards.

We examine how risk diversity can affect a population's ability to achieve a target threshold to effectively avoid a disastrous outcome. From there, we explore the benefits that introducing financial incentives can have on overall welfare, cooperation and target achievement. As our main contributions, we show that:

(1) Symmetrical diversity in risk leads to asymmetrical changes in agents policies, reducing overall contributions, target achievement and welfare;
(2) Higher risk values reduce the consequences of risk diversity;
(3) Higher target achievement is not always equivalent to higher welfare;
(4) Risk diversity offers opportunities in the design of financial incentives, which can foster higher cooperation, target achievement and global welfare than the ones obtained in the absence of diversity.

We begin our paper with Section 2 on related work. In Section 3, we model the collective risk dilemma, explain how risk diversity is introduced, describe the agents' learning dynamics, and present the structure of the designed financial incentives. Finally, we display our results in Section 4 and conclude our paper in Section 5.

2 RELATED WORK

In this section, we delve into the literature on risk and diversity in collective games, we outline the ways in which risk diversities can appear in a society, and present usages of financial incentives to promote cooperation in social dilemmas.

Previous works on CRDs, both experimental and theoretical, have concluded that a higher risk translates into a higher probability that agents cooperate and, consequently, might help in escaping the tragedy of the commons [6, 11, 30, 31, 40, 42].

The global risk factor is not the only decisive factor in an agent's willingness to cooperate. The introduction of different inequalities between agents was shown to have a significant impact on cooperation. Under evolutionary game theory, wealth and returns inequality were found to reduce cooperation in a continuous public goods game [23]. Similar results were found for wealth inequalities in CRDs, in experimental settings [51], as well as under evolutionary game theory [57], and reinforcement learning [30].

While most sources of heterogeneity studied in the literature focus on wealth inequality, we argue that risk — either perceived or effective — is yet another source of heterogeneity worth studying in populations facing collective risk dilemmas. A survey of 119 countries confirmed significant variance in public concern and risk assessment of the global climate change problem [27]. The Global Risk Report reveals how participants assess differently the risks of different global concerns [18]. Similarly, the recent COVID-19 pandemic led to the distinction between people at normal risk and those at increased risk of severe illness from COVID-19. The World Health Organization identified medical conditions that can increase the risk of getting seriously ill [34]. Other governmental units have classified jobs into different risk exposure levels [33], and several other studies and reports have been published highlighting similar diversities [5, 17, 35, 53]. We note that risk diversity can either emerge from a subjective diversity in risk perception or an objective diversity in actual risk exposure. In this work, we focus on the consequences of diversity in risk exposure.

In games with collective risks, FI were studied as a means to prevent free-riding and promote cooperation among agents, both in homogeneous and symmetrical games [21, 48], as well as in heterogeneous games with wealth inequality [13].

FI in the form of peer-rewards or institutional rewards can also promote cooperation in public goods games [8, 14, 37, 38, 49, 59].

In the context of reinforcement learning, the introduction of decaying peer-rewards during training, was empirically proven to be effective for avoiding the tragedy of the commons in a resource appropriation setting [28] and increase the chances of converging to a pro-social equilibrium in a Stag-Hunt game [58].

The literature on rewards and punishments in social dilemmas, recognizes financial incentives as an instrument to prohibit free-riding [46]. In our work we identify a novel purpose for FI emerging from the challenges raised by diversity in risk exposure. Under such diversity, an agent at low risk of facing a disaster can choose to defect, not to free-ride on the contributions of others, but simply because of an indifference towards target achievement. For an agent at low risk, if the threshold is not met, the chances of facing a disaster are low. If the cost of cooperating is equal or higher than the cost of failing, an agent can lose interest in target achievement. For instance, vaccinated individuals with low probability of facing severe illness from COVID-19 can feel discouraged to follow unpleasant safety measures. Nevertheless, a vaccinated person can still be contagious, and the selfish behavior can have dangerous consequences for other people at high risk. From a climate action perspective, people in different locations can be at different risks of facing extreme weather, forest fires, floods etc. and therefore can experience different levels of motivation to achieve climate agreements. We investigate how financial incentives can provide a new solution for aligning interests among heterogeneous agents, rather than only a way to discourage free-riding.

In this section, we highlighted from the literature the effect that the risk and different heterogeneities have on cooperation in collective risk dilemmas. We identified a lack of papers investigating risk heterogeneities despite their prevalence in real-life interactions. In this paper, we address this gap and investigate the consequences that risk diversity has on populations facing collective risks. We exploit our findings to propose a new use case for financial incentives in non-symmetrical social dilemmas.

3 MODEL

Dynamics in a multi-agent system depend both on the decision making dynamics of the agents composing it and on the environment dynamics where interactions take place. In this section, we detail the dynamics of the collective risk dilemma that agents are engaged in, the learning dynamics of the individual agents, as well as the dynamics of the added financial incentives.

3.1 Game Dynamics

A collective risk dilemma (CRD) is a game in which agents need to cooperate to avoid an eventual disaster [11, 31, 40, 44, 56, 57]. Agents' success in avoiding the disaster requires a minimum collective effort. Effort is modeled by the costly contribution of players towards a common pool. If contributions are below the threshold they will not alleviate the consequences of the disaster. Additionally, all contributions above the threshold do not create any additional value for the players. As a result, agents are on the one hand motivated to cooperate to ensure that the disaster is avoided, and on the other hand motivated to defect and free-ride hoping that others will put in the required efforts for disaster avoidance.

Formally, in a population of finite size Z, we allocate for every player an initial endowment b. Players are then sampled into groups of size N to play CRDs. They need to jointly collect enough contributions to reach a target threshold $\mathbf{t}$ to avoid with certainty a common disaster. If a group manages to achieve the threshold target, the disaster is avoided and players only lose what they had contributed to the common pool. However, should the target not be met, agents, depending on their level of risk exposure to the disaster, will lose a fraction p of their remaining endowment. That is, every agent i, will incur additional disaster losses with a risk probability r_i. At the end of the game, player i who started with an initial endowment b will be left with

$$b^i_{final} = \begin{cases} (1-c_i)b & \text{if the disaster was avoided,} \\ (1-c_i)b - p(1-c_i)b & \text{otherwise.} \end{cases} \tag{1}$$

where c_i represents the binary choice of either contributing 0 or a fraction c of the initial endowment to the pool ($c_i \in \{0, c\}$).

The perceived benefits or harm of these losses in endowment is a subjective function known as the *utility* in economic game theory. One common utility function is the log-utility. The log-utility function has been used when studying the impact of wealth inequality in CRDs [30] and is used more broadly in economy to capture what is known as a diminishing marginal utility [36]. It assumes that losing money is perceived as more painful by poorer individuals than by richer ones. Similarly, it assumes that losing $x\%$ of one's wealth is equally painful for anyone, regardless of how much wealth that percentage represents in absolute value. While all agents are equally wealthy in our scenario, we do intend to examine mixtures of heterogeneities in future works, such as the combination of wealth inequality with risk diversity. With that in mind, to better compare our results with future works, we decide to adopt a log-utility function. Moreover, in several scenarios, results under log-utility do not differ from those under linear utility [30]. The payoffs of the game are expressed as the difference in the log of agents' wealth before and after a game was played. Avoiding a disaster will cost a cooperator $x_C = \log\left(\frac{b-cb}{b}\right) = \log(1-c)$, and a defector $x_D = \log\left(\frac{b}{b}\right) = 0$ or nothing. Facing a disaster will cost cooperators $\bar{x}_C = \log(1-c-p(1-c))$ and defectors $\bar{x}_D = \log(1-p)$. The goal of each player is to find a probabilistic strategy π^*_i — representing the probability of player i choosing to cooperate — that maximizes the expected payoff.

3.1.1 Introduction of risk diversity. We consider risk diversity in the form of binary risk *classes*. That is, we split our population into two classes: agents at high risk of being affected by the disaster and agents at low risk. Agents at high risk represent a fraction z_H of the population, and agents at low risk represent the remaining fraction $z_L = 1 - z_H$ of the population. Given an average population risk value r and a risk diversity value δ, if the target is not achieved, agents at high risk will lose an additional fraction p of their remaining wealth with probability $r_H = r + \frac{1}{2z_H}\delta$ while agents at low risk only face that disaster with a risk probability $r_L = r - \frac{1}{2z_L}\delta$ — this guarantees that the average risk in the population remains r regardless δ.

3.1.2 Numerical values. The population size is set to $Z = 200$ individuals, representative of the number of countries engaged in the climate action problem, a commonly studied CRD. The agents are organized in groups of $N = 6$. They are given an initial endowment $b = 1$ and can choose to either cooperate and contribute a fraction $c = 0.1$ of it to a common pool or defect and contribute nothing. Participants have stochastic policies π_i and will sample for each game one of the two choices according to the probabilities defined by their policies. The threshold $\mathbf{t}$ is set so that the target is only achieved if at least half of the agents in a group cooperate, i.e., $\mathbf{t} = Mcb$ with $M = \frac{N}{2}$. Agents at high and low risk are equally frequent in the population with $z_H = z_L = 50\%$ of the population. This means, for an average population risk r and a risk diversity δ, agents at high risk will face a disaster with probability $r_H = r + \delta$ while agents at low risk will face a disaster with a risk probability $r_L = r - \delta$. If the threshold target is not achieved, every agent that faces a disaster pays a penalty of $p = 0.7$ or 70% of its remaining wealth. We proceed with two experiments: in the first, we investigate the effect of increasing mean risk values r, while in the second, we explore the consequences of increasing risk diversities δ. This allows us to better understand the impact of risk diversity for regimes of high and low baseline risk (δ fixed and varying r) and also the impacts of risk diversity in the form of symmetric risk distribution (r fixed and varying δ).

3.2 Agent Learning Dynamics

The goal of the paper is to understand how simple reinforcement dynamics can encourage or discourage cooperative behaviors in populations with risk diversity. The findings of this study can later be used by works that aim to engineer cooperative capabilities for agents in these scenarios. We choose to model the agents learning dynamics using the Roth-Erev Algorithm which was shown to successfully model human decision making in social dilemmas [39].

Accordingly, we create a population of Z agents and allow every player i, to assign and update a propensity value for each of the possible actions. In the 2-actions collective risk dilemma, this translates into a propensity vector $\mathbf{q}_{i,k} = \left[q_{i,k}(C), q_{i,k}(D)\right]^T$ where $q_{i,k}(C)$ and $q_{i,k}(D)$ are the respective propensities for the cooperative and the defective action after the k^{th} learning interaction. The resulting propensities are normalized using the soft-max function to derive the stochastic policy $\pi_{i,k}$, used to sample an action at the next interaction. At the end of the interaction, when returns are distributed, every player i, depending on the selected action A and the received reward x, updates its propensity vector such that

$$\begin{aligned} q_{i,k+1}(A) &= (1-\phi)q_{i,k}(A) + x \\ q_{i,k+1}(\neg A) &= (1-\phi)q_{i,k}(\neg A) \end{aligned} \tag{2}$$

where $\neg A$ represents the non-chosen action and ϕ is a forgetting parameter that inhibits the propensities from growing to infinity.

We adopt the same population training algorithm used previously to train a population of agents facing a CRD under wealth inequality [30]. We summarize the procedure in Algorithm 1. At every step k, a group of N agents is selected randomly from the population of Z agents. The agents in this group engage in a CRD. Every player j in the group chooses randomly one of the 2 available actions following probabilities $\pi_{j,k-1}$, derived by normalizing the propensity vector $\mathbf{q}_{j,k-1}$. The selected actions and the agents' risk factors determine their different payoffs. The payoffs are then used by the agents in the group to update their propensity vectors. This is repeated for a total of K learning-steps. Since the algorithm does not guarantee that all agents are chosen equally as many times, we define K', the minimum number of learning-steps per agent, and ensure that training continues until every agent performs at least K' learning-steps. We choose the same numerical values used by Merhej et al. and fix $K = 2.5 \times 10^6$, $K' = 3 \times 10^4$, and $\phi = 1 \times 10^{-3}$. We repeat all simulations for 5 runs.

Algorithm 1: Training a population of independent RL agents using the Roth-Erev algorithm

Init: K *total number of learning-steps*, K' *minimum number of learning-steps per agent*
for $i \leftarrow 1$ **to** Z, *population size* **do**
  $\mathbf{q}_{i,0} \leftarrow$ *random initialization*;
  $u_i \leftarrow 0$ `/* number of updates per agent */`
for $k \leftarrow 1$ **to** K **do**
  1. *sample random group* G *of size* N;
  2. *sample actions* $A_j \sim \pi_{j,k-1}$ *for* $j \in G$;
  3. *evaluate disaster avoidance for* $j \in G$;
  4. *calculate* log-utility *payoff of* $j \in G$;
  5. *update* $\mathbf{q}_{j,k}$ *(Eq. 2)*;
  6. $u_j \leftarrow u_j + 1$ *for* $j \in G$;
  7. $u_{min} \leftarrow min(\mathbf{u})$
while $u_{min} < K'$ **do**
  repeat steps 1. to 7.

3.3 Financial Incentives

Achieving cooperation in a multi-player collective risk social dilemma is challenging. The introduction of risk diversity among participants augments the problem with additional difficulties. While agents at high risk of facing a disaster are highly motivated to achieve the target, agents at a lower risk, may not feel the same urge. The indifference of agents at low risk makes the target harder to achieve for agents at high risk. Because the game itself offers little incentive for agents at low risk to cooperate, we investigate how a financial motivation from agents at high risk can drive disinterested agents to cooperate. We explore financial incentives (FI) as a zero-sum transfer of a reward from agents at high risk to agents at low risk. The peer-rewards should align the motives of agents at different risk levels and ensure cooperation among agents with originally different interests.

In this paper we do not focus on how FI among peers can emerge in a self-organized way but rather study the benefits of the existence of a pre-arranged agreement between the two classes. We detail below the design of the FI and the conditions for a transfer to occur.

In a group of N participants engaging in a CRD, let N_H be the number of agents at high risk, and $N_L = N - N_H$, the number of agents at low risk. We use the superscript to indicate cooperative and defective agents. That is, out of N_L agents at low risk, we have N_L^C cooperators and $N_L^D = N_L - N_L^C$ defectors. The same notation can be applied to agents at high risk. After an interaction, if the target threshold is achieved, agents at high risk will transfer a fraction of their wealth to agents at low risk as an incentive to motivate cooperative actions. To outbalance cooperation costs, every agent at high risk will contribute a fraction $f = 2c$ of the original endowment, i.e., twice the cooperation cost, to a pool dedicated for financial incentives. The collected sum ($2cb \times N_H$) by agents at high risk is then equally distributed among N_L^C cooperators at low risk.

Agents at high risk who started with an initial endowment b, will end the game with

$$b^H_{final} = \begin{cases} b - cb - 2cb & \text{if they cooperated,} \\ b - 2cb & \text{if they defected.} \end{cases} \tag{3}$$

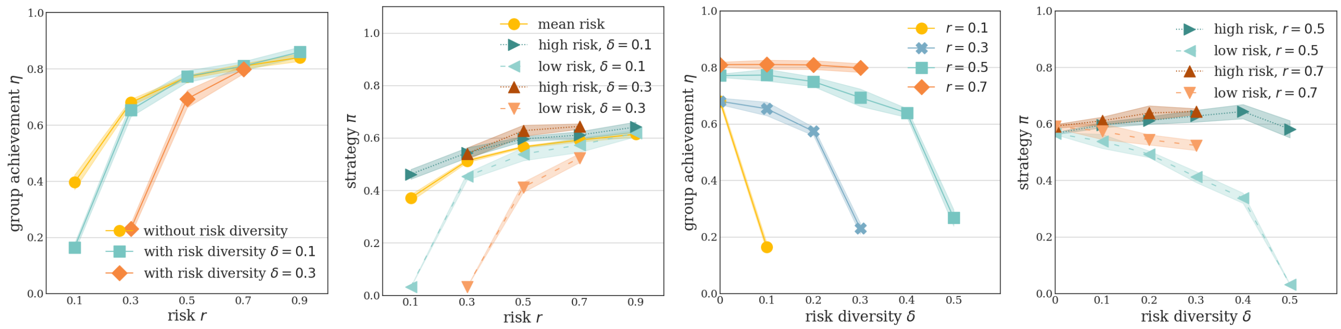

(a) (b) (c) (d)

Figure 1: (a) Group achievement of populations with and without risk diversity w.r.t. the risk r. In all cases, group achievement increases with r. (b) Learned strategies of agents from a population without diversity (mean risk) and those of agents from populations with risk diversities $\delta = 0.1$ and $\delta = 0.3$ w.r.t. the risk r. With higher risk, the relative impact of the diversity decreases and agents behave more like homogeneous populations. (c) Group achievement of populations under different mean risk w.r.t. the risk diversity δ. Risk diversity decreases overall group achievement despite a constant average risk. (d) Learned strategies of agents at high and low risk in populations of different mean risk w.r.t. the risk diversity δ. For high diversity, the reduction in cooperation from agents at low risk is not compensated by an equal increase from agents at high risk and overall contributions decrease. In all the panels, shaded areas represent the standard deviation over 5 runs.

On the other hand, agents at low risk who started with an initial endowment b, will finish the game with

$$b^L_{final} = \begin{cases} b - cb + 2cb\frac{N_H}{N_L^C} & \text{if they cooperated,} \\ b & \text{if they defected.} \end{cases} \tag{4}$$

Again, agents receive as payoffs the difference in the log of their wealth before and after an interaction.

If the target threshold is not achieved, no rewards are transferred and agents receive the payoffs of the original game of Section 3.1.

4 RESULTS

We study the consequences of risk diversity for populations of agents learning to play the game introduced in section 3.1 with the RL algorithm of section 3.2. After the learning phase, the strategies are evaluated based on a) the resulting population's probability of achieving the target threshold $\mathbf{t}$, and b) the amount of remaining wealth in that population. To get those values, for every setting, we rollout a game where the population is split into groups of N players. In each group, agents, following their learned strategies, choose to either contribute or not. At the end of the game, we calculate η, the percentage of groups in the population that successfully reach the target, as well as ζ, the fraction of secured welfare or remaining wealth in the population after cooperation costs and disaster losses. The two variables are evaluated and averaged over 10^6 simulations. Studies are run both on heterogeneous populations with risk diversity, as well as on their homogeneous counterparts (i.e., populations with the same average risk r but no diversity δ).

4.1 Effect of Risk Inequality on Cooperation and Target Achievement

To study the effect that risk inequality can have on a population facing a CRD, we begin by comparing the group achievement rate η and the learned strategies of a homogeneous population on one hand, with those of heterogeneous populations with risk diversity factors $\delta = 0.1$ and $\delta = 0.3$ on the other hand. We plot the results for varying average risk factors r in Figures 1a and 1b.

With or without risk diversity, we observe that the group achievement rate increases with the risk (Figure 1a). This comes from the higher willingness of agents to cooperate (Figure 1b) as the costs of failure increase with the risk. The results we show are consistent with other studies on collective risks [30, 40].

However, while most studies report that inequalities have a decisive impact on group achievement [30, 41, 57], the impact of risk diversity is mitigated at higher risk values. The consequences of risk diversity are largest for $\delta = r$, when the relative strength of the diversity δ/r is the highest. This value decreases as r increases, resulting in achievements similar to homogeneous populations.

The observation remains true for the learned strategies. In fact, as r increases, the strategies learned by both classes seem to converge to the strategies learned by a homogeneous population.

Upon our first findings, we investigate the role of the diversity factor δ. In a second experiment, we examine the variations in target achievement and cooperation with respect to the risk diversity δ. In Figure 1c, we observe how for the same average risk, stronger diversity causes a drop in achievement. Again, the effect is mitigated for larger risk values ($r = 0.7$) and amplified at $r = \delta$, when all the population's risk is only carried by half of the population, i.e. where $r_L = r - \delta = 0$ and $r_H = r + \delta = 1$. The conclusions are coherent with the results in Figure 1a.

Figure 1d shows the strategies followed by individuals at high and at low risk in populations with mean risk $r = 0.5$ and $r = 0.7$. Results for $r = 0.1$ and $r = 0.3$ are qualitatively similar to the results for $r = 0.5$ and are omitted for more visibility. For stronger diversity, we notice an increased gap in cooperation rates between the two classes as agents at low risk cooperate less and agents at high risk compensate by cooperating more. However, as the relative strength of the diversity r/δ approaches 1, i.e., $\delta = r$, the reduction in cooperation of agents at low risk is not compensated by a similar increase in cooperation from agents at high risk. As a consequence, total contributions in the population decrease which explains the drop in target achievement. When risk diversity is relatively strong, and especially when it concentrates responsibility of target achievement on a subset of the population, target achievement and overall contributions of a population significantly decline. We notice that this does not occur for $r = 0.7$ since values of $\delta > 0.3$ are not supported.

4.1.1 Reinforcement learning dynamics and Nash equilibria. We highlight the distinctions between solutions found under learning with reinforcement dynamics and the game theoretical Nash equilibria. The general Nash equilibrium is a point where no agent can increase personal payoff by deviating *alone* from the chosen strategy profile. The Nash equilibrium considers fully independent players with no pre-established coordination between them. While it is difficult to extract Nash equilibria for n-player threshold games [10], we can prove that total defection is at least one of the possible Nash solutions. In the case where $Z - 1$ agents in the population defect, the best response of the Z^{th} agent is to defect as well. This is because no agent can achieve the target alone and cooperation would only result in unnecessary costs. When learning under reinforcement dynamics, whether with or without risk diversity, no population converged to the defective Nash equilibrium. We confirmed that this result remains true for simulations where agents start with a stronger initial preference for defection.

Additionally, we can check if the learned strategies are Nash equilibria. For all tested risk values, if we fix $Z - 1$ strategies in the population and only allow the Z^{th} agent to change its strategy, defection is the most profitable choice. In that sense, learned strategies are not Nash equilibria. Populations of reinforcement learners do not converge to defective or other possible Nash solutions. We hypothesize that the large size of the population can hamper convergence to Nash equilibria for adaptive agents. After a costly cooperative act, if several agents in a population simultaneously increase their defection rate, the next interaction may become less profitable as the increase in failure (caused by a reduction in target achievement) is not compensated by the individual decrease in cooperation cost. Assessing the benefits of diverging alone from a strategy profile, which is necessary information for computing Nash equilibria, is not easily done in populations of reinforcement learners where several agents simultaneously change their strategies. Learning with reinforcement dynamics in large populations seems to help in escaping defective Nash equilibria.

4.2 Effect of Risk Inequality on Secured Population Welfare

While we have shown that target achievement increases with the risk r and decreases with the diversity δ, we question how this translates to the total welfare of the population. We plot respectively in Figures 2a and 2b, the fraction of remaining welfare for a) fixed diversities δ and varying risks r, and b) for a fixed average risk $r = 0.5$ and varying risk diversities δ. Losses in welfare can either come from cooperation costs or disaster occurrences.

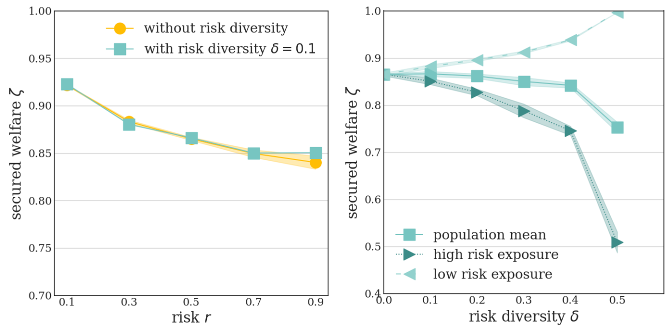

(a) (b)

Figure 2: (a) Total population welfare when following learned strategies for a fixed $\delta = 0.1$ and varying risk factors r. Total welfare decreases with the risk despite higher target achievement rates. (b) Population welfare distribution between low at risk and high at risk agents for $r = 0.5$ and varying diversity factors δ. Total welfare remains relatively constant but drops for $\delta = 0.5$ when the losses of agents at high risk are not compensated by the gains of agents at low risk.

Despite the fact that target achievement increases with the risk (Figure 1a), Figure 2a indicates that the total welfare of a population drops with higher risk values. Although agents increase their contributions and achieve the target threshold more often, the increase is not enough to compensate for the elevated risk of a disaster. Higher risk values are often considered a means for escaping the tragedy of the commons [31, 40, 42]. Our findings suggest that this can be true if a disaster's real risk is smaller than its perceived risk by the agents. This is because a real increase in risk is not met with a sufficient increase in cooperation and target achievement to cover the larger losses from higher risk exposure. Higher target achievement is not equivalent to higher welfare.

Similarly, while target achievement decreased with risk diversity, Figure 2b suggests that the tendency is not directly translated to the total welfare. As δ increases and η decreases, agents at high risk are more exposed to disasters and suffer larger losses. However, simultaneously, agents at low risk are less exposed to disasters and save on cooperation costs. For $\delta \leq 0.4$, the total losses and gains in the population balance out and the average welfare remains relatively unchanged. Lower target achievements do not strictly imply larger losses.

Nonetheless, this result does not extend to all risk diversity values. For $\delta = 0.5$, although risk diversity remains symmetrical in the population, the resulting gains and losses in welfare in each class are non-symmetrical. The asymmetry is such that the losses of agents at high risk are greater than the gains of agents at low risk. The population as a whole incurs additional losses from risk diversity. The findings demonstrate how in the case of strong diversity, ignoring heterogeneities and averaging out can lead to over-optimistic results and an underestimation of the expected losses.

In the next section, we investigate how financial incentives can help a population of reinforcement learners reduce its losses despite high risk diversity.

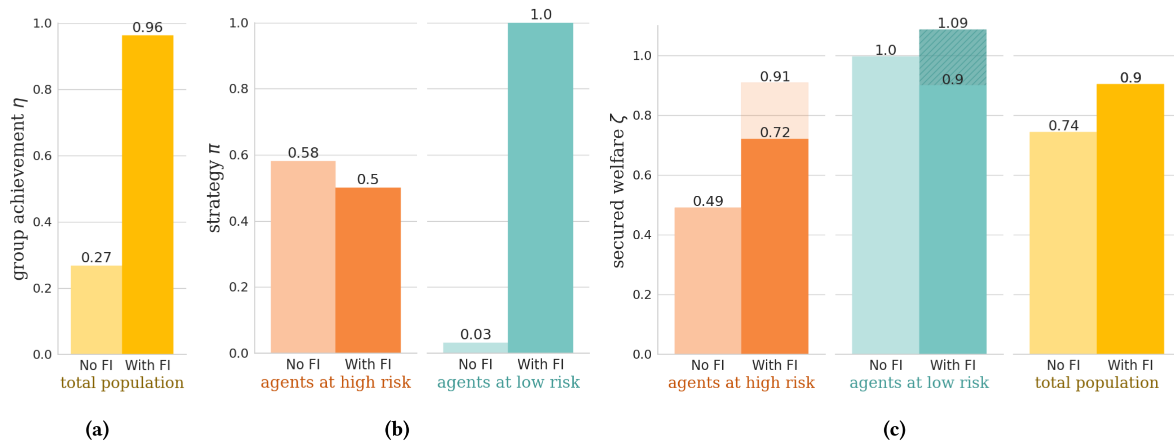

Figure 3: (a) Overall group achievement for populations engaging in a CRD of average risk $r = 0.5$ and risk diversity $\delta = 0.5$ with and without financial incentives (FI). We observe a clear rise in target achievement with FI. (b) The corresponding cooperation rates learned by agents at high and low risk in each case. FI positively encourage agents at low risk to cooperate. This allows agents at high risk to cooperate a little less and save on cooperation costs. (c) Welfare of agents at high and low risk as well as the total welfare of the population in situations with and without FI. In the case of FI, we plot for agents at high and low risk, both the secured wealth before and after reward transfer. Thanks to higher cooperation from agents at low risk, agents at high risk face disasters less often and secure more of their initial welfare (from 49 to 91%). This is only possible if they sacrifice a fraction of that wealth (in light color) and transfer it to agents at low risk. Meanwhile, agents at low risk who now cooperate more, incur cooperation costs and their welfare drops (from 100 to 90%). This is compensated by the new received incentives from agents at high risk (hatched dark color). Both agents transferring financial incentives and agents receiving them in exchange for cooperative contributions manage to increase their wealth. FI create a new win-win equilibrium.

4.3 Financial Incentives

We uncovered how risk diversity can decrease overall contributions in a population, reduce target achievement, and in some cases result in unnecessary losses in resources. Here, we take the extreme case of $r = \delta = 0.5$ where cooperation, threshold reaching and total welfare are at their minima. We examine how financial incentives (FI) can improve on each of the three criteria. We train a population of RL agents of average risk $r = 0.5$ and risk diversity $\delta = 0.5$, to play a CRD with additional FI as described in Section 3.3. In this setting, agents at low risk face no risk at all ($r_L = r - \delta = 0$). Defection from agents at low risk is not motivated by free-riding but by an indifference towards target achievement.

We observe in Figure 3a that FI can significantly increase group achievement from 27% to 96%. The almost certainty with which the population meets the target, cannot be achieved by a unilateral cooperation of agents at high risk. Figure 3b plots the cooperation levels of agents at high and low risk in settings with and without FI. While agents at low risk never cooperate in the original setting, FI actively motivate them to do so. We observe a switch from total defection to total cooperation. Meanwhile, this allows agents at high risk to reduce their cooperation from 58% to around 50% and save on cooperation costs.

While FI increase group achievement, they are costly actions for agents at high risk. We have seen earlier that higher target achievement does not necessarily translate into higher welfare. In Figure 3c we plot the average welfare of each class of agents, with and without FI. Through increased target achievement and avoidance of disasters, agents at high risk increase their secured welfare from $\zeta = 49\%$ to 91%. Yet, the increased wealth is only possible if agents at high risk reward their peers at low risk. Almost half of the gained wealth is spent on FI which eventually reduces the total welfare to $\zeta = 72\%$. Nevertheless, despite the costs of financial incentives, the final secured wealth of agents at high risk remains higher than the secured wealth without FI.

Next, we examine the consequences of FI on agents at low risk. Without FI, agents at low risk, to minimize cooperation costs, almost never cooperate. Since they are at no risk of facing disasters, their wealth remains untouched ($\zeta = 1$). With FI, agents consistently cooperate and lose on average $c = 10\%$ of their wealth decreasing ζ to 90%. However, this is supplemented by incentives from agents at high risk which in turn increases their welfare to new highs of 109% of their initial wealth.

From a population's perspective, FI are internal exchanges of rewards. Although no wealth is actively produced, financial incentives create new pro-social equilibria, leading to overall increased gains in welfare (from 74 to 90%).

Our results confirm the powerful advantages of zero-sum FI in a population facing collective risks under risk diversity. FI can mitigate and overcome the challenges imposed by risk diversity such as the decrease in target achievement, cooperation and secured welfare. Moreover, FI not only recover from risk diversity, but also improve on settings without diversity. Figures 1c and 1d show that populations without risk diversity ($\delta = 0$), achieve the target threshold with probability $\eta = 80\%$ and an average cooperation $\pi = 60\%$. Financial incentives enhance these performances and achieve the target at a rate of $\eta = 96\%$ and an average cooperation of $\pi = 75\%$.

While risk diversity decreases cooperation and target achievement, one can profit from risk heterogeneity to design incentive schemes which foster levels of collective success that are even higher than those obtained in the absence of heterogeneity. Yet, the collective benefits do not completely abolish the internal inequalities in the population. Considerable differences in cooperation and secured wealth still exist between the two classes. This raises the challenge of designing incentives that can attain even higher levels of fairness and equality within a population.

5 CONCLUSIONS

We examine how risk inequality between RL agents can affect a population's target achievement rate and the cooperation levels of different risk classes. First, we find that high risk diversity (large values of δ/r) causes a noticeable reduction in group achievement. As diversity increases, cooperation levels of agents at high and low risk respectively increase and decrease. However, while the change in risk exposure is symmetrical between the classes, the changes in cooperation are not. The asymmetry is always such that the increase in cooperation of one class is smaller than the accompanied decrease in cooperation of the other class. This raises significant target achievement difficulties. Despite challenges in cooperation, populations of reinforcement learners never converge to fully defective Nash equilibria.

Next, we find that changes in target achievement cannot be directly mapped to changes in welfare. We show that welfare can decrease despite increased target achievement or remain constant despite a reduction in target achievement. Depending on the application, it can be relevant to consider the target achievement or the total welfare as a measure of the effectiveness of a strategy, or even search for conciliating solutions in between. Nevertheless, for high risk diversities, we observe again an asymmetry in welfare gains and losses among the two classes. Agents at high risk lose more welfare because of disaster occurrences than agents at low risk gain from reduced disaster exposure and cooperation costs.

Finally, we propose to leverage risk heterogeneity using financial incentives that can simultaneously increase target achievement, global contributions and the welfare of both agents at high and at low risk. Moreover, the achieved performances are higher than the ones obtained in homogeneous populations. However, we note that while financial incentives can raise global performance of heterogeneous populations above those of similar homogeneous populations, they do not fully eliminate inequalities in the population. Designing FI that promote higher fairness and equality in a population can be an interesting topic for future works.

We note that, in this paper, we studied diversity in risk exposure assuming that agents recognize their true risk of exposure to a disaster. In reality, the real risk of exposure to a disaster is often not known and people behave according to how likely they assess or perceive an imminent danger to be. Based on the risk perception, people plan mitigation and adaptation policies to respectively mitigate or reduce the likelihood of a disaster occurring (here, increasing η) and adapt or reduce the impact of the possible damages in the case of disaster occurrence (here, decreasing the cost of failure p) [32]. Investigating cooperation in the context of collective risks is analogous to evaluating mitigation policies. Cooperation in the form of contributions helps a group of agents to achieve a target threshold and avoid a possible disaster.

We have shown in our study, that when agents are at different risk exposure levels (e.g., healthy and unhealthy people in the face of pandemics, people in geographic locations that are more or less prone to extreme weather etc.), effective mitigation policies are harder to implement, i.e., the target achievement drops. This is because the population at low risk is reluctant to cooperate, making the problem harder to solve for the rest of the population at high risk. Here, we have shown that financial incentives from agents at high to agents at low risk can overcome the emerging difficulties. When the usage of such agreements is not possible, our results demonstrate the urgent implementation of proper adaptation measures (building flood defenses, increasing hospital capacity etc.), as the likelihood of a disaster increases with increased risk diversity.

Additionally, we found that agents do not sufficiently increase their contributions with increased risk. Higher target achievements are achieved, but they do not compensate for the increased losses of ever more likely disasters. Mitigation policies, although effective, must still be accompanied by adaptation policies to alleviate the damages of possible disasters. Because agents contribute more but not sufficiently more with the risk, another conceivable solution would be to increase risk *perception* of participants. Convincing players that they are at a higher risk than they actually are, can incite them to contribute more and avoid disaster losses.

The most alarming setting happens when agents perceive a risk lower than the real risk. We know that agents do not contribute enough with higher risks. A further reduction in contributions because of a false sense of security can be highly dangerous. This is the case of the climate change problem in underdeveloped countries [27]. Education may be crucial for aligning the risk perceptions of agents with the actual risk of the problem and achieving long term cooperation. Nevertheless, education and knowledge propagate slowly and through several generations. In the face of an imminent danger, financial incentives again, from countries perceiving the high danger to countries not assessing the emergency, can be a fast and practical way of increasing global contributions and mitigating the problem.

In the context of developing cooperative AI, our findings under risk diversity suggest that cooperation cannot be enforced using communication, retaliation or other classical solutions for symmetrical social dilemmas. Mixing individualistic and social qualities in agents is necessary to achieve cooperative AI under diversity. An illustrative example is the idea of hyper-rational choices in which actors think about profit or loss of other actors in addition to their personal profits [2]. For fully individualistic agents, allowing inter-agent contracts, bargains and financial incentives, can be a way for selfish cooperation to emerge. This requires the understanding of the payoffs of the game, the capacity to develop win-win proposals, and the ability to implement those contracts (i.e, the ability to receive and provide rewards from and to other agents).

ACKNOWLEDGMENTS

This work has received funding from the EU's H2020 program (grant 76595) and FCT-Portugal (grants UIDB/50021/2020, PTDC/MAT-APL/6804/2020, and PTDC/CCI-INF/7366/2020).

REFERENCES

[1] Fahad Alanezi, Arwa Althumairi, et al. 2021. A comparative study of strategies for containing the COVID-19 pandemic in Gulf Cooperation Council countries and the European Union. *Inform. Med. Unlocked* 23 (2021), 100547.
[2] Gholamreza Askari, Madjid Eshaghi Gordji, and Choonkil Park. 2019. The behavioral model and game theory. *Palgrave Commun.* 5, 1 (2019), 1–8.
[3] Robert Axelrod. 1980. Effective choice in the prisoner's dilemma. *J. Conflict. Resolut.* 24, 1 (1980), 3–25.
[4] Robert Axelrod and William Donald Hamilton. 1981. The evolution of cooperation. *Science* 211, 4489 (1981), 1390–1396.
[5] Emily A Benfer and Lindsay F Wiley. 2020. Health justice strategies to combat COVID-19: protecting vulnerable communities during a pandemic. *Health Aff Blog* 10 (2020).
[6] Manuel Chica, Juan M Hernández, and Jacques Bulchand-Gidumal. 2021. A collective risk dilemma for tourism restrictions under the COVID-19 context. *Sci. Rep.* 11, 1 (2021), 1–12.
[7] Manuel Chica, Juan M Hernández, and Francisco C Santos. 2022. Cooperation dynamics under pandemic risks and heterogeneous economic interdependence. *Chaos Solit.* 155 (2022), 111655.
[8] Ross Cressman, Jie-Wen Song, Bo-Yu Zhang, and Yi Tao. 2012. Cooperation and evolutionary dynamics in the public goods game with institutional incentives. *J. Theor. Biol.* 299 (2012), 144–151.
[9] Astrid Dannenberg, Andreas Löschel, Gabriele Paolacci, Christiane Reif, and Alessandro Tavoni. 2011. Coordination under threshold uncertainty in a public goods game. *ZEW-Centre for European Econ Research Disc Paper* 11-065 (2011).
[10] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. 2009. The complexity of computing a Nash equilibrium. *SIAM J. Comput.* 39, 1 (2009).
[11] Elias F. Domingos, Jelena Grujić, Juan C Burguillo, Georg Kirchsteiger, Francisco C Santos, and Tom Lenaerts. 2020. Timing uncertainty in collective risk dilemmas encourages group reciprocation and polarization. *iScience* 23, 12 (2020), 101752.
[12] Elias F. Domingos, Jelena Grujić, Juan C Burguillo, Francisco C Santos, and Tom Lenaerts. 2021. Modeling behavioral experiments on uncertainty and cooperation with population-based reinforcement learning. *Simul Model Pract Theory* 109, 102299 (2021).
[13] Yali Dong, Shuangmei Ma, Boyu Zhang, Wen-Xu Wang, and Jorge M Pacheco. 2021. Financial incentives to poor countries promote net emissions reductions in multilateral climate agreements. *One Earth* 4, 8 (2021), 1141–1149.
[14] Yali Dong, Boyu Zhang, and Yi Tao. 2016. The dynamics of human behavior in the public goods game with institutional incentives. *Sci. Rep.* 6, 1 (2016), 1–7.
[15] Ido Erev and Alvin E Roth. 2014. Maximization, learning, and economic behavior. *Proc Natl Acad Sci USA* 111 (2014), 10818–10825.
[16] Jakob Foerster, Richard Y Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. 2018. Learning with opponent-learning awareness. In *Proc. of AAMAS'18.* 122–130.
[17] Organisation for Economic Co-operation and Development. 2020. *COVID-19: Protecting people and societies.* https://www.oecd.org/coronavirus/policy-responses/covid-19-protecting-people-and-societies-e5c9de1a/ (accessed: 26.08.2021).
[18] World Economic Forum. 2021. *The Global Risk Report 2021.* https://www.weforum.org/reports/the-global-risks-report-2021 (accessed: 22.09.2021).
[19] Drew Fudenberg and David K Levine. 1998. *The theory of learning in games.* MIT Press.
[20] Elizabeth Gibney. 2020. Whose coronavirus strategy worked best? Scientists hunt most effective policies. *Nature* 581, 7806 (2020), 15–17.
[21] António R Góis, Fernando P Santos, Jorge M Pacheco, and Francisco C Santos. 2019. Reward and punishment in climate change dilemmas. *Sci. Rep.* 9, 1 (2019).
[22] 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.
[23] Oliver P Hauser, Christian Hilbe, Krishnendu Chatterjee, and Martin A Nowak. 2019. Social dilemmas among unequals. *Nature* 572, 7770 (2019), 524–527.
[24] Alexis Jacq, Julien Perolat, Matthieu Geist, and Olivier Pietquin. 2019. Foolproof Cooperative Learning. *ArXiv:1906.09831* (2019).
[25] Peter Kollock. 1998. Social dilemmas: The anatomy of cooperation. *Annu. Rev. Sociol.* 24, 1 (1998), 183–214.
[26] Veronica M Lamarche. 2020. Socially connected and COVID-19 prepared: The influence of sociorelational safety on perceived importance of COVID-19 precautions and trust in government responses. *Soc. Psychol. Bull* (2020).
[27] Tien Ming Lee, Ezra M Markowitz, Peter D Howe, Chia-Ying Ko, and Anthony A Leiserowitz. 2015. Predictors of public climate change awareness and risk perception around the world. *Nat. Clim. Chang.* 5, 11 (2015), 1014–1020.
[28] Andrei Lupu and Doina Precup. 2020. Gifting in multi-agent reinforcement learning. In *Proc. of AAMAS'20.* 789–797.
[29] Michael W Macy and Andreas Flache. 2002. Learning dynamics in social dilemmas. *Proc Natl Acad Sci USA* 99 (2002), 7229–7236.
[30] Ramona Merhej, Fernando P Santos, Francisco S Melo, and Francisco C Santos. 2021. Cooperation between Independent Reinforcement Learners under Wealth Inequality and Collective Risks. In *Proc. of AAMAS'21.* 898–906.
[31] Manfred Milinski, Ralf D Sommerfeld, Hans-Jürgen Krambeck, Floyd A Reed, and Jochem Marotzke. 2008. The collective-risk social dilemma and the prevention of simulated dangerous climate change. *Proc Natl Acad Sci USA* 105, 7 (2008).
[32] NASA. 2021. *Climate Change Adaptation and Mitigation.* https://climate.nasa.gov/solutions/adaptation-mitigation/ (accessed: 06.10.2021).
[33] OSHA (U.S. Department of Labor). 2020. *Hazard Recognition.* https://www.osha.gov/coronavirus/hazards (accessed: 26.08.2021).
[34] World Health Organization. 2020. *Coronavirus disease (COVID-19): Risks and safety for older people.* "https://www.who.int/news-room/q-a-detail/coronavirus-disease-covid-19-risks-and-safety-for-older-people" (accessed: 26.08.2021).
[35] Tasnime Osama, Bharat Pankhania, and Azeem Majeed. 2020. Protecting older people from COVID-19: should the United Kingdom start at age 60? *J R Soc Med* 113, 5 (2020), 169–170.
[36] Ole Peters and Murray Gell-Mann. 2016. Evaluating gambles using dynamics. *Chaos* 26, 2 (2016), 023103.
[37] Flávio L Pinheiro and Fernando P Santos. 2018. Local Wealth Redistribution Promotes Cooperation in Multiagent Systems. In *Proc. of AAMAS'18.* 786–794.
[38] David G Rand, Anna Dreber, Tore Ellingsen, Drew Fudenberg, and Martin A Nowak. 2009. Positive interactions promote public cooperation. *Science* 325, 5945 (2009), 1272–1275.
[39] Alvin E Roth and Ido Erev. 1995. Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games Econ. Behav.* 8, 1 (1995), 164–212.
[40] Francisco C Santos and Jorge M Pacheco. 2011. Risk of collective failure provides an escape from the tragedy of the commons. *Proc Natl Acad Sci USA* 108, 26 (2011), 10421–10425.
[41] Francisco C Santos, Marta D Santos, and Jorge M Pacheco. 2008. Social diversity promotes the emergence of cooperation in public goods games. *Nature* 454, 7201 (2008), 213.
[42] Francisco C Santos, Vitor V Vasconcelos, Marta D Santos, PNB Neves, and Jorge M Pacheco. 2012. Evolutionary dynamics of climate change under collective-risk dilemmas. *Math. Models Methods Appl. Sci.* 22 (2012), 1140004.
[43] Fernando P Santos, Simon A Levin, and Vítor V Vasconcelos. 2021. Biased perceptions explain collective action deadlocks and suggest new mechanisms to prompt cooperation. *iScience* 24, 4 (2021), 102375.
[44] Fernando P Santos, Samuel Mascarenhas, Francisco C Santos, Filipa Correia, Samuel Gomes, and Ana Paiva. 2020. Picky losers and carefree winners prevail in collective risk dilemmas with partner selection. *Auton Agent Multi Agent Syst* 34, 2 (2020), 1–29.
[45] Fernando P Santos, Jorge M Pacheco, Francisco C Santos, and Simon A Levin. 2021. Dynamics of informal risk sharing in collective index insurance. *Nature Sustainability* (2021), 1–7.
[46] Karl Sigmund. 2010. *The calculus of selfishness.* Princeton University Press.
[47] Brian Skyrms. 2010. *Signals: Evolution, learning, and information.* Oxford University Press.
[48] Weiwei Sun, Linjie Liu, Xiaojie Chen, Attila Szolnoki, and Vítor V Vasconcelos. 2021. Combination of institutional incentives for cooperative governance of risky commons. *iScience* 24, 8 (2021), 102844.
[49] Attila Szolnoki and Matjaz Perc. 2010. Reward and cooperation in the spatial public goods game. *EPL (Europhysics Letters)* 92, 3 (2010), 38003.
[50] Attila Szolnoki and Matjaž Perc. 2012. Evolutionary advantages of adaptive rewarding. *New J. Phys.* 14, 9 (2012), 093016.
[51] Alessandro Tavoni, Astrid Dannenberg, Giorgos Kallis, and Andreas L"oschel. 2011. Inequality, communication, and the avoidance of disastrous climate change in a public goods game. *Proc Natl Acad Sci USA* 108, 29 (2011), 11825–11829.
[52] Pham Tien Thanh et al. 2020. Survey data on government risk communication and citizen compliance during the COVID-19 pandemic in Vietnam. *Data in brief* 33 (2020), 106348.
[53] Kristin van Barneveld, Michael Quinlan, Peter Kriesler, Anne Junor, Fran Baum, Anis Chowdhury, Pramod N Junankar, Stephen Clibborn, Frances Flanagan, Chris F Wright, et al. 2020. The COVID-19 pandemic: Lessons on building more equal and sustainable societies. *Econ. Labour Relat. Rev.* 31, 2 (2020), 133–157.
[54] Ibo Van de Poel and Jessica Nihlén Fahlquist. 2013. Risk and responsibility. In *Essentials of risk theory.* Springer, 107–143.
[55] Ibo Van de Poel, Jessica Nihlén Fahlquist, Neelke Doorn, Sjoerd Zwart, and Lamber Royakkers. 2012. The problem of many hands: Climate change as an example. *Sci. Eng. Ethics* 18, 1 (2012), 49–67.
[56] Vitor V Vasconcelos, Francisco C Santos, and Jorge M Pacheco. 2013. A bottom-up institutional approach to cooperative governance of risky commons. *Nat. Clim. Chang.* 3, 9 (2013), 797–801.
[57] Vítor V Vasconcelos, Francisco C Santos, Jorge M Pacheco, and Simon A Levin. 2014. Climate policies under wealth inequality. *Proc Natl Acad Sci USA* 111, 6 (2014), 2212–2216.
[58] Woodrow Z Wang, Mark Beliaev, Erdem Bıyık, Daniel A Lazar, Ramtin Pedarsani, and Dorsa Sadigh. 2021. Emergent Prosociality in Multi-Agent Games Through Gifting. *arXiv preprint arXiv:2105.06593* (2021).
[59] Jia-Jia Wu, Cong Li, Bo-Yu Zhang, Ross Cressman, and Yi Tao. 2014. The role of institutional incentives and the exemplar in promoting cooperation. *Sci. Rep.* 4, 1 (2014), 1–6.