

Multi-agent Reinforcement Learning in the All-or-Nothing Public Goods game on Networks

Benedikt Valentin Meylahn
University of Amsterdam
Amsterdam, The Netherlands
b.v.meylahn@uva.nl

ABSTRACT

We study interpersonal trust by means of the all-or-nothing public goods game between agents on a network. The agents are endowed with the simple yet adaptive learning rule, *exponential moving average*, by which they estimate the behavior of their neighbors in the network. Theoretically we show that in the long-time limit this multi-agent reinforcement learning process always eventually results in indefinite contribution to the public good or indefinite defection (no agent contributing to the public good). However, by simulation of the pre-limit behavior, we see that on complex network structures there may be mixed states in which the process seems to stabilize before actual convergence to states in which agent beliefs and actions are all the same. In these metastable states the local network characteristics can determine whether agents have high or low trust in their neighbors. More generally it is found that more dense networks result in lower rates of contribution to the public good. This has implications for how one can spread global contribution toward a public good by enabling smaller local interactions.

KEYWORDS

Multi-agent learning; Social dilemma; Trust; Public goods

ACM Reference Format:

Benedikt Valentin Meylahn. 2025. Multi-agent Reinforcement Learning in the All-or-Nothing Public Goods game on Networks. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, Detroit, Michigan, USA, May 19 – 23, 2025, IFAAMAS, 9 pages.

1 INTRODUCTION

Interpersonal trust is important to the healthy functioning of societies. When trust is in abundance this enables interaction in the absence of costly contracts, monitoring and policing. In this paper we study interpersonal trust as it may be learned by agents in a group setting. To do this we use the public goods game [3, 5, 10, 14, 27, 31–35] in the context of agents who learn based on experience. We study the impact of group size and the effect of complex network structure.

The public goods game consists of a group of players who must decide whether or not pay for a public good. Traditionally, the benefit to all players is proportional to the number of players who decide to contribute (or cooperate). Contribution to the public good

however, comes at a cost. A selfish agent maximizes their reward by not contributing as long as the others do. In this paper we consider an all-or-nothing version in which the benefit is only rewarded when all players pay the cost.

We are motivated by situations in real life in which each participant has a veto power. Consider for instance the sharing of a water source, if all involved handle it with care and avoid pollution it remains useful. However, if even one firm pollutes the source, this ruins the situation for all others. Indeed, a threshold public goods game is well suited to the modeling of such an ecological conundrum [3, 10, 34]. Another example is slightly reversed though still applicable: Choosing a competitive or a collusive price, can be conceived of as a public goods game. If all players set the collusive price, the ‘public good’ is attained (from the perspective of the firms, for society this is not good at all). However, if but one firm sets the competitive price, the party is ruined for everyone else as they reap the profits of the greatest market share.

In the following sections we provide a non-exhaustive review of the literature on the public goods game. We focus on: all-or-nothing versions, evolutionary approaches, spatial (networked) versions, and experience based agent learning.

1.1 Literature: General

Evolutionary dynamics are a prominent tool for scholars studying the evolution of cooperation in social dilemmas. The public goods game is no exception as evidenced by the amount of work in which this is the case (see for example [14, 27, 31–33, 35]). In the evolutionary setting players get rewards for their performance in the game in the form of fitness. Players with a greater fitness are more likely to pass their strategy on to progeny. In these studies long-term stable strategies are identified by means of Monte Carlo simulation or by reducing the system to the replicator equation.

In the spatial version of the public goods games players are modeled as vertices on a graph and the interactions are defined by connections in the graph. In particular it is common to draw the group of agents for a game by considering the neighborhood of a randomly drawn player. Common network topologies used in the study of the spatial evolutionary public goods game include the torus [14, 16, 36] and grids [32, 33]. There has also been work done on evolutionary dynamics of the public goods game on complex networks [7, 12, 27].

The effect of group size on the possibility of cooperation in the public goods game is an ongoing debate with findings in both empirical and theoretical lines of research going both ways. The interested reader is referred to [2, 13, 18] for a positive influence of group size on cooperation, and [23, 37] for negligible and sometimes even negative effect in empirical research). In theoretical



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

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

research, Suzuki and Akiyama [31], for example, show that in the context of evolutionary dynamics well-mixed population including a reputation mechanism, the level of cooperation decreases as the group size increases. In contrast, Szolnoki and Perc [33] show that on grids, larger groups lead to more cooperation up to an optimal group size, after which the level of cooperation decreases again.

1.2 Literature: All-or-nothing, and learning

Contribution by all players may be required to avoid disaster, and so we believe the all-or-nothing version of the public goods game to be an important object of study. However, there is not an abundance of literature on this topic. Studying agents that learn and adapt their behavior based on experience is relatively well established for fully-mixed populations (see for example [20, 28–30] or [15] for two-player experience based learning). Yet applying experience based learning to agents on a network is still a growing field (see for instance [1, 5]). As such we now present some the work that relates to these two topics: all-or-nothing public goods, and experience based learning.

Mailath and Postlewaite [17], and more recently Bierbrauer and Winkelmann [4] study the all-or-nothing version of the public goods game from a mechanism design perspective. Mailath and Postlewaite [17] show that for a single public good in which all players have a veto power, as the group size grows, the probability of realizing the public good goes to zero. Bierbrauer and Winkelmann [4] on the other hand show that if players instead vote on a bundle of public goods, the probability of realizing the public good stays positive dependent on the capacity (size of the bundle).

Szolnoki and Perc [32] consider a version of the game which includes the all-or-nothing version as a special case. In particular they study the effect of a ‘critical mass’ (minimum level of contribution from the group as a whole) required for the benefit to be realised in the spatial evolutionary public goods game on a grid. They find that there is an optimal critical mass yielding a maximum level of cooperation.

Bladon and Galla [5] study the dynamics of the public goods game with agents (on a ring-and-hub graph) who use batch learning. They find that increasing the degree of the hub player (connected to all others on the ring) increases their propensity to defect. There are studies that consider experience based learning for agents in the coordination game [1, 20, 21]. In particular the agents in [1, 21] use Q-learning to coordinate opinions with their neighbors on random geometric graphs. Meanwhile in [20] agents on a fully connected graph use exponential moving average with a constant learning rate to coordinate in pairwise interactions. In both of these learning methodologies the long-term dynamics exhibit convergence to one strategy (full coordination).

1.3 Research gap and contributions

The advancement of computing technology makes it possible to study a population of sophisticated agents who learn based on past experiences in games. By the nature of social interactions we are motivated to study these on networks of ‘interesting’ topologies. As such in this paper we aim to address this research gap by studying the all-or-nothing public goods game with experience based agent learning for a population on complex networks.

We contribute to multi-agent learning on networks by our theoretical finding that any connected network of agents using the exponential moving average learning rule eventually converges to a single pure strategy. We do this without a separation of time scales which is often used to remove the stochasticity and thereby simplify the analysis (and may result in other dynamics, see [6, 28–30]). This contributes to understanding dynamics of interacting agents who learn in network environments beyond the calculation of the expected dynamics.

We also contribute to the understanding of stability of contribution in the public goods game on a network. In particular we show how network properties such as average degree, and network density effect which pure strategy the dynamics converge to. Additionally, we show that on the intermediate time scale, the dynamics may exhibit metastable states in which local network properties strongly influence the behavior of the agents. This highlights the negative effect of interaction size on the likelihood of contributing also in the context of multi-agent learning from experience. We also illustrate that this metastable behavior emerges with network heterogeneity by comparing the dynamics on the random geometric network to those on regular networks (which exhibit fast convergence to steady state).

Our results suggest that to foster contribution on a large scale, it may help to identify local groups where cooperation can yield intermediate benefits. Threading these local groups together without increasing the size of the interaction then might allow this cooperation to spread.

2 THE MODEL

First we introduce the base game as if it were a one shot game. Then we define more formally the repeated and population version.

2.1 Base game

The all-or-nothing public goods game consists of $k \in \mathbb{N}$ players acting simultaneously. These players choose whether to contribute (C) or defect (D). Before making their choice each player i observes a private random variable $\lambda_i \in \mathbb{R}$ which is the payoff they would receive if they and all other players contribute to the public good. For convenience we define the generic random variable such that each λ_i is distributed as the generic $\lambda \in \mathbb{R}$.

The players are not aware of the actions of other players until after all decisions have been made, and they are also not privy to the possible reward obtained by other players (λ_j for $j \neq i$). Players that defect (play D) obtain a payoff of 1 regardless of the actions of other players. If all players contribute (C), then each player i for $i \in \{1, \dots, k\}$ obtains the reward λ_i . Players that contribute while at least one player did not contribute, obtain reward 0.

2.2 Population model of the all-or-nothing public goods game

We model a population of $N \in \mathbb{N}$ agents who interact in groups K_t of size $k_t \in \mathbb{N}_{\leq N}$ at time steps $t \in \mathbb{N}$. At each time step each agent i in the group K_t observes a random variable¹ (drawn iid)

¹The idea of each agent observing a randomly drawn utility from the awarded public good is not common yet also not novel. In particular a similar approach is taken in [4, 17] for the context of mechanism design. In the context of population dynamics

$\lambda_i(t) \in \mathbb{R}$ which determines the size of their gain if the public good is achieved. The agent starts with an endowment of 1 and is required to choose whether or not to contribute 1 to the public good. For each agent $i \in K_t$ we define their action at time t :

$$A_i(t) = \begin{cases} C, & \text{if they contribute} \\ D, & \text{if they defect.} \end{cases} \quad (1)$$

Subsequently the agent payoff $\pi_c(t)$ defined for taking the contribute action in round $t \in \mathbb{N}$ is²:

$$\pi_c(t) = \begin{cases} \lambda_i(t), & \text{if } \sum_{j \in K_t} \mathbb{1}_{\{A_j(t)=C\}} = k, \\ 0, & \text{else.} \end{cases} \quad (2)$$

The payoff for taking the defect action is identically one: $\pi_d(t) = 1$.

2.3 Agent belief

Each agent $i \in [1, \dots, N]$ holds a belief on the probability that a random other agent will contribute to the public good $x_i(t)$ (initialized independently for each agent $x_i(0) \sim \mathcal{U}[0, 1], \forall i$). For those agents $i \in K_t$ drawn to play the game at time t the belief update follows exponential moving average with fixed learning step size α :

$$x_i(t+1) = x_i(t)(1-\alpha) + \alpha \left(\frac{\sum_{j \in K_t \setminus i} \mathbb{1}_{\{A_j(t)=C\}}}{k_t - 1} \right). \quad (3)$$

Agents who were not chosen to interact simply retain their most recent belief. The exponential moving average rule has been studied in the context of two-player games in a series of papers authored by Sato, Akiyama, Farmer, and Crutchfield in various configurations [28–30]. These studies however, make use of a separation of time scales under which the random dynamics converge to the expected dynamics. We study the random dynamics when this separation of time scales does not hold as done for the stochastic coordination game of [20].

2.4 Agent actions

The agents don't know the distribution of the reward but only observe the realization of their random variable $\lambda_i(t)$ which they receive if they and all other agents in the round contribute. They combine this with their belief on which actions are likely to be taken by others. This yields the expected utility per action at time t to agent j ($u_i^c(t)$ for cooperating and $u_i^d(t)$ for defecting):

$$u_i^c(t) = \lambda_i(t)x_i(t)^{k_t-1} \quad (4)$$

$$u_i^d(t) = 1. \quad (5)$$

We assume that the agents act myopically. That is they chose the action which maximizes their one-round expected utility. Thus, agent i takes the contribute action in round t if $u_i^c(t) \geq u_i^d(t)$:

$$x_i(t)^{k_t-1} \lambda_i(t) \geq 1. \quad (6)$$

and repeated interactions [25, 26, 39] use a random payoff matrix observed by each player in the game drawn anew each round.

²Awarding the public good only if all agents take the contribute action has been studied for one-shot games in [4, 17]. The all-or-nothing version we consider is a special case of the threshold version studied experimentally in [10, 34] and theoretically in [32, 35]. These studies however do not consider experience based agent learning or complex network structures as we do.

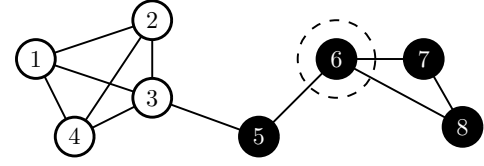


Figure 1: Illustration of how the players are selected to play a round of the all-or-nothing public goods game on a graph of 8 players. The focal player (encircled by a dashed line) is $i = 6$. The players taking part in the game are thus $K_t = N[6] = \{6, 5, 7, 8\}$ (indicated by filled nodes) and so $k_t = 4$.

We use the weak inequality for convenience. Because $\lambda_i(t)$ is continuous, a tie occurs with probability zero, and so this is not an issue. This can be rearranged to see that the agent contributes if $x_i(t)^{k_t-1} \geq 1/\lambda_i(t)$. For convenience we define the cdf of the random variable $1/\lambda$ derived from the generic random variable λ as F :

$$F(x) := \mathbb{P}(1/\lambda \leq x), \quad (7)$$

and observe that for an outside observer (who does not know the realization of $\lambda_i(t)$ but does know $x_i(t)$) the probability of an agent contributing follows $F(x^{k_t-1})$. Note that F is not associated with any t because the distribution of $\lambda_i(t)$ is always that of λ for all $i \in K_t$ and all $t \in \mathbb{R}$.

An intuitive explanation of the model is provided in Appendix A (of the arXiv version: arxiv:2412.20116).

3 PUBLIC GOODS ON NETWORKS

Now we present our results for the stochastic all-or-nothing public goods game with multi-agent learning on networks.

We define a graph $G = (V, E)$ with vertices representing agents $|V| = N$ and the edges $(u, v) \in E$ representing connections between them. In each discrete round indexed by $t \in \mathbb{N}$, a vertex $i(t) \in V$ (called the ‘focal’ player) is selected uniformly at random from V . To be precise $i(t) \sim \mathcal{U}\{1, N\}$ for all $t \in \mathbb{N}$ (there is no correlation between the vertex drawn at round t and $t+1$)³. Subsequently the selected player’s closed neighborhood $N[i(t)] = \{v : (i(t), v) \in E\} \cup \{i(t)\}$ plays a round of the all-or-nothing public goods game described in §2. This means that the game in round $t \in \mathbb{N}$ with agent $i(t)$ as the focal player has interaction size $k_t = |N[i(t)]|$ and players $K_t = N[i(t)]$. The way in which agents are selected to take part in a round of the all-or-nothing public goods game on a network is illustrated in Figure 1.

3.1 Main theoretical result

Now we state the main result which guides our simulations of the model. The result states that as time goes to infinity, the belief and actions of all agents converge to the same action belief pair; either having full trust ($x \rightarrow 1$) and always contributing or having no trust ($x \rightarrow 0$) and always defecting. This extends a similar result for a fully-mixed (no network) population and pairwise interactions [20].

PROPOSITION 3.1. *For learning rate $\alpha \in (0, 1)$ and some $\epsilon \in (0, \alpha/d)$ where $d > 0$ is the maximum degree in the graph minus*

³Introducing a correlation between vertices drawn as the focal vertex, as long as there is positive probability on each vertex to be drawn in each round, does not influence the main result.

one⁴, if $F(0) = 0$, $F(1) = 1$, $F(x) \in (0, 1)$ for all $x \in (0, 1)$ and there exist finite order derivatives of $F(x)$ at $x = 0$ and $x = 1$ which are finite, then the probability of all agents converging in belief and action asymptotically is one, i.e.

$$\mathbb{P}(\exists t_0 : \{\mathbf{x}(t) \in [0, \epsilon]^N, \forall t \geq t_0\}, \text{ or } \{\mathbf{x}(t) \in [1 - \epsilon, 1]^N, \forall t \geq t_0\}) = 1. \quad (8)$$

This proposition means that there can be no disagreement on trust in the very long term. Eventually all agents agree that everyone is to be trusted, or that none is to be trusted.

The proof of this proposition depends on two intermediate results. The first of which states that absorption in the corners is possible:

LEMMA 3.2 (ABSORPTION IN THE CORNERS IS POSSIBLE). *Let $\alpha \in (0, 1)$, $\epsilon \in (0, \alpha/d)$, where $d > 0$ is the maximum degree in the graph minus one and suppose $\mathbf{x}(t_0) \in [0, \epsilon]^N$ for some $t_0 \in \mathbb{N}$, then*

$$\mathbb{P}(\mathbf{x}(t) \in [0, \epsilon]^N, \forall t \geq t_0 \mid \mathbf{x}(t_0) \in [0, \epsilon]^N) > 0. \quad (9)$$

For similar α and ϵ , if $\mathbf{x}(t_0) \in [1 - \epsilon, 1]^N$ for some $t_0 \in \mathbb{N}$, then

$$\mathbb{P}(\mathbf{x}(t) \in [1 - \epsilon, 1]^N, \forall t \geq t_0 \mid \mathbf{x}(t_0) \in [1 - \epsilon, 1]^N) > 0. \quad (10)$$

The proof closely follows the proof of Lemma 1 in Meylahn et al. [20]. We sketch the proof for absorption in $x \in [0, \epsilon]^N$ with attention to the changes due to the network structure and group interaction.

Proof sketch. In order to absorb in the corner $[0, \epsilon]^N$ from round $t_0 \in \mathbb{N}$ we require all agents to play defect for all rounds $t \geq t_0$. We define the probability of the event of all agents playing defect in rounds $t_0, t_0 + 1, \dots, t_0 + n$ given that all agents start with belief $x \in [0, \epsilon]$ as D_n .

Observe first that any finite connected graph may be covered by cliques which share connections. As such we proceed by induction on the cliques of the network. As base case we consider a graph consisting of one clique of size $k \geq 2$, and call $D_n(k)$ the probability of all k agents playing defect in all rounds $t_0, t_0 + 1, \dots, t_0 + n$, given that all agents started with belief $x \in [0, \epsilon]$.

Now for the base case $D_n(k)$ is bounded:

$$D_n(k) \geq \prod_{i=0}^n (1 - F(z_i^{k-1}))^k, \quad (11)$$

where $z_0 = \max_j \{x_j(t_0)\}$, and $z_i = (1 - \alpha)z_{i-1}$ for all $i \geq 1$. Subsequently we take the logarithm of both sides and the limit as $n \rightarrow \infty$:

$$\lim_{n \rightarrow \infty} \log(D_n) \geq \lim_{n \rightarrow \infty} k \sum_{i=0}^n \log(1 - F(z_i^{k-1})). \quad (12)$$

We observe that $1 - 1/y \leq \log(y)$ for all $y \in \mathbb{R}_{>0}$, which we use by setting $1 - F(z_i^{k-1}) = y$ to show:

$$\log(1 - F(z_i^{k-1})) \geq \frac{F(z_i^{k-1})}{F(z_i^{k-1}) - 1}. \quad (13)$$

Applying the above to (12) we get

$$\lim_{n \rightarrow \infty} \log(D_n(k)) \geq k \lim_{n \rightarrow \infty} \sum_{i=0}^n \frac{F(z_i^{k-1})}{F(z_i^{k-1}) - 1}. \quad (14)$$

⁴Note that for any connected graph with $N > 2$ the maximum degree is bigger than 1.

Proceeding with similar arguments to the proof of Lemma 1 in [20],⁵ we observe that the right hand side of (14) converges to a negative number. Taking the exponential of both sides gives the desired result that $\lim_{n \rightarrow \infty} D_n(k) > 0$.

As induction step we assume that L cliques of size $k_i \geq 2$ for $i \in \{1, \dots, L\}$ absorb in the corner at $[0, \epsilon]^N$ with positive probability and call this probability $b_L > 0$.

Once we have shown that $b_{L+1} > 0$ then our proof is complete. We construct the induction step graph by referring to its L ‘original’ cliques and the added clique number $L + 1$. We remind the reader that the new graph is still connected, thus there are agents in clique $L + 1$ who share edges with agents in the L original cliques.

Because we require all agents to play defect in *all* rounds from t_0 to $t_0 + n$ as $n \rightarrow \infty$ we know that each agent will be selected an infinite number of times as the focal agent. Because we are dealing with multiplication we can rearrange the terms from an arrangement by ascending rounds to infinite products of rounds with specific agents as the focal agent. We denote by k_i the number of agents in the game in round $t_0 + i$ for $i \in \{0, 1, \dots, n\}$. We define new notation for legibility:

$$\mathbb{P}_{[a,b]^k}(\cdot) := \mathbb{P}(\cdot \mid x(t_0) \in [a, b]^k). \quad (15)$$

Thus by definition of D_n we have:

$$D_n = \prod_{i=0}^n \mathbb{P}_{[0,\epsilon]^N}(A(t_0 + i) = D \mid A(t) = D, \forall t_0 \leq t < t_0 + i). \quad (16)$$

We collect the rounds in which the focal agent is *a*) the L original cliques and *not* connected also to an agent in clique $L + 1$ in set M_a , *b*) the clique $L + 1$ and *not* connected to an agent in the L original cliques in set M_b , and *c*) one of the agents who is connected to both clique $L + 1$ and the L original cliques in set M_c . So we rewrite the above product:

$$\begin{aligned} D_n &= \prod_{t \in M_a} \underbrace{\mathbb{P}_{[0,\epsilon]^N}(A(t) = D \mid A(\tau) = D, \forall t_0 \leq \tau < t)}_{=:a} \\ &\times \prod_{t \in M_b} \underbrace{\mathbb{P}_{[0,\epsilon]^N}(A(t) = D \mid A(\tau) = D, \forall t_0 \leq \tau < t)}_{=:b} \\ &\times \prod_{t \in M_c} \underbrace{\mathbb{P}_{[0,\epsilon]^N}(A(t) = D \mid A(\tau) = D, \forall t_0 \leq \tau < t)}_{=:c}. \end{aligned} \quad (17)$$

We can split the rounds in M_c further by each agent which, when taking the limit as $n \rightarrow \infty$, results in separate infinite products per agent. These are positive by reasoning similar to the base case and noting that each agent’s belief is the same or lower than in the base case because they have also played other games (where all agents played D by the conditioning) and so $c > 0$.

In the same way $b > 0$ (and $a > 0$) because this follows only the agents in clique $L + 1$ (the original graph) whose per round belief is thus bounded in the same way as in the base case (or by the induction assumption), with the agents on the boundary who

⁵This makes use of Abel’s convergence test for this product of sequences, and the ratio test with L’Hôpital’s rule to show convergence. In this step we require F to have a finite order derivative with a finite value.

have played in more games which by the conditioning resulted in all agents playing D and thus had even lower beliefs resulting in the defect action at a greater probability. Thus $a \times b \times c > 0$. \square

The other intermediate result states that from a non-corner, the corners may be reached at positive probability. Define I_N as the space of the agent belief excluding the corners, $I_N := [0, 1]^N \setminus ([0, \epsilon]^N \cup [1 - \epsilon, 1]^N)$.

LEMMA 3.3 (REACHING THE CORNERS IS POSSIBLE). *Let $\alpha \in (0, 1)$ and $\mathbf{x}(t_0) \in I_N$ for some $t_0 \in \mathbb{N}$ and $\epsilon \in (0, \alpha/d)$, where $d > 0$ is the maximum degree in the graph minus one, then*

$$\mathbb{P}_{I_N}(\exists t_1 \in \mathbb{N} : \mathbf{x}(t_1) \in [0, \epsilon]^N, t_1 > t_0) > 0,$$

and

$$\mathbb{P}_{I_N}(\exists t_1 \in \mathbb{N} : \mathbf{x}(t_1) \in [1 - \epsilon, 1]^N, t_1 > t_0) > 0,$$

The proof of this lemma closely follows the proof of Lemma 3 in Meylahn *et al.* [20]. Here we sketch the differences due to network effects for reaching $\mathbf{x} \in [0, \epsilon]^N$. The procedure for proving it for $\mathbf{x} \in [1 - \epsilon, 1]^N$ is similar.

Proof sketch. We proceed in two cases. In the first case it follows similarly to the proof of Lemma 3 in [20] that the population may reach a corner, while in second case we show that it is possible to reach the starting point of the first case completing the proof.

Case 1. $\mathbf{x}(t_0) \in [0, 1 - \epsilon]^N$: We first delineate a finite path of positive probability by which a population of N agents reaches the corner given that each agents belief starts in $(0, 1 - \epsilon)$.

An agent's belief reaches $x \leq \epsilon$, at worst after κ rounds with all agents playing D in each round where

$$\kappa = \left\lceil \frac{\log(\epsilon)}{\log(1 - \alpha)} \right\rceil. \quad (18)$$

For α and ϵ as defined, this is a finite number. For agents with belief $x \in (0, 1 - \epsilon)$ the probability of playing defect in a round is positive.

Consider a vertex cover of the graph and label the agents in it arbitrarily $j = 1, \dots, v$. Suppose we have t_0 such that the population belief is $\mathbf{x}(t_0) \in [0, 1 - \epsilon]^N$ and in rounds $t_0 + i$ for $i = 0, 1, \dots, \kappa v$ the focal player $j = i \bmod v$ is chosen and all agents always play defect in each round. This selection of players happens at probability $(1/N)^{\kappa v} > 0$. The probability of all agents always playing defect is also positive. This is because their belief $x(t_0) \leq 1 - \epsilon$, and so $(1 - F(x_{t_0}^{k_0-1}))^{k_0} > 0$, starts positive, and increases each time they have played. They play a finite number of games and thus this finite product of positive numbers converges to a positive number.

Case 2. $\mathbf{x}(t_0) \notin [0, 1 - \epsilon]^N$ Now suppose instead that $\mathbf{x}(t_0) \notin [0, 1 - \epsilon]$, implying that at least one agent has belief $x(t_0) > 1 - \epsilon$. Suppose as worst case scenario that all agents have belief $x = 1$ except agent $j = 1$ with belief $x_1 = 1 - \epsilon$. Now we show that with positive probability this population can reach the state $\mathbf{x} \in [0, 1 - \epsilon]$. Let the degree of agent 1 be n and note that $1 - \alpha + \alpha(\frac{n-1}{n}) \leq 1 - \epsilon$, by the choice of ϵ .

Agent 1 is chosen as focal player two rounds in a row at probability $1/N^2$. If in the first of those two rounds, agent 1 plays defect and all others play contribute, the beliefs of all other agents in the game becomes $x \leq 1 - \epsilon$, while agent 1's belief becomes $x > 1 - \epsilon$. In the second round however the other n agents play defect at positive probability, and so after both rounds, at positive probability p , agent 1 and all of their n neighbors have beliefs $x \leq 1 - \epsilon$.

Continuing in this way along a spanning tree of the graph, at each iteration playing two games (the first to bring new players into belief $x \leq 1 - \epsilon$, and the second to ensure that the focal player's belief is also $x \leq 1 - \epsilon$) each agent in the population has belief $x \leq 1 - \epsilon$. The largest number of rounds required is for the path graph on N players which results in $2(N - 1)$ games. Thus at probability greater than or equal to $1/N^{2(N-1)} p^{N-1} > 0$ in at most $2(N - 1)$ rounds, all agents have belief $x \leq 1 - \epsilon$. From here the proof continues as in case 1. \square

Together these results prove Proposition 3.1. To see this note that after $r < \infty$ rounds, the beliefs of the agents may enter one of the corners at positive probability (by Lemma 3.3) and be absorbed there (by Lemma 3.2). Call the probability of this entering and absorbing event $p_a > 0$. Then the probability of never absorbing in one of the two corners behaves like $\lim_{n \rightarrow \infty} (1 - p_a)^n = 0$.

The assumptions on F are minor, and reasonable; if an agent believes there is zero (conversely 100%) chance of others contributing then they should take the defect (conversely contribute) action with probability one as this surely maximizes their expected reward.

3.2 Simulation setup

Because Proposition 3.1 only guarantees convergence as $t \rightarrow \infty$, by simulation we study the effect of the network on the speed of convergence to steady state and on the pre-limit behavior. We fix the learning rate $\alpha = 0.3$ in favor of faster simulation runs. Similarly to balance having interesting network topologies and a manageable computational load we fix the population size $N = 50$. For this numerical simulation we also fix $F(x) = x^{1/4}$.

We set an $\epsilon_s = 10^{-4}$, in combination with a maximum number of rounds $T = 10^7$. We stop each iteration if either all agents have belief $x < \epsilon_s$ or $x > 1 - \epsilon_s$ as proxy for convergence or, alternatively if the maximum number of rounds is reached.

To study the effect of network topology we use the random geometric network model [9, 24]. This random network model is relevant because it shares properties of actual social networks [19]. The graph is constructed by fixing a radius $r_g \in (0, 1)$ and placing N points uniformly at random in the unit square $[0, 1]^2$. Subsequently, any two points that are at distance $d < r_g$ are connected by an edge. In particular we focus on connected networks and so for each iteration we draw random geometric graphs repeatedly until we draw a connected geometric graph. For the settings described we run 500 iterations for each setting $r_g \in \{0.15, 0.2, 0.25, 0.3\}$.

3.3 Random geometric network results

We plot the time to convergence $\tau := \min\{t : \mathbf{x}(t) \in [0, \epsilon_s]^N \text{ or } \mathbf{x}(t) \in (1 - \epsilon_s, 1]^N\}$ in Figure 2 in the form of tail probabilities⁶ ($\mathbb{P}(\tau \geq t)$). We also tabulate the end state of the simulation runs in Table 1. In Figure 2 we notice that increasing the network radius (r_g) increases the probability of quicker convergence. Notice that for $r_g = 0.15, 0.2$ the majority of the runs do not reach consensus in the simulated time which may also be read off Table 1. By the linear shape of this tail probability plotted on a log-log scale we

⁶Note that this decision is informed by Proposition 3.1. If we did not know that convergence happens eventually with probability one, we might falsely plot the probability of convergence rather than the probability of convergence *before* $t \in \mathbb{N}$.

r_g	0.15	0.20	0.25	0.3
Contribution	0.48	0	0	0
Defection	0	0.366	0.948	0.994
Not converge	0.52	0.634	0.052	0.006

Table 1: Portion of runs out of 500 which converged to contribution and defection as well as the proportion which did not converge in 10^7 rounds.

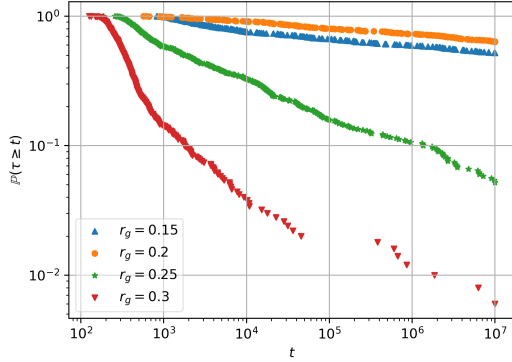


Figure 2: Tail probabilities ($\mathbb{P}(\tau \geq t)$) on a log-log scale for the time to convergence for different values of the radius r_g used in the random geometric graph model to sample networks.

posit that a heavy-tailed distribution underlies the time to consensus on a random geometric graph. To check whether this is really the case we turn to the broadest sub-class of heavy-tailed distributions: subexponential distributions (examples are the Pareto, and LogNormal distributions, for an introduction see [22]). There is equivalence between distributions satisfying the catastrophe principle and subexponential distributions. A distribution H is said to satisfy the catastrophe principle if for $n \geq 2$

$$\frac{\mathbb{P}(\max\{X_1, \dots, X_n\} > t)}{\mathbb{P}(X_1 + \dots + X_n > t)} \rightarrow 1, \quad (19)$$

as $t \rightarrow \infty$ where X_i for $i = 1, \dots, n$ are iid and have distribution H . The intuitive understanding of this principle is that a large sum of these random variables is likely caused by a single massive variable⁷. In Table 2 we present an estimate of the ratio of (19) for the time to consensus. In particular the ratio is shown for $t = 10^6$, and $n = 2$. To estimate each ratio we sample pairs of simulations 250 times once with and once without replacement of samples.

r_g	0.15	0.20	0.25	0.3
NR:	0.848:0.852	0.896:0.896	0.196:0.196	0.02:0.02
R:	0.824:0.824	0.9:0.9	0.192:0.192	0.028:0.028

Table 2: Ratio between $\mathbb{P}(\max\{X_1, X_2\} > t) : \mathbb{P}(X_1 + X_2 > t)$ where the X s are our simulated time to convergence. NR: no replacement (run i paired with run $i + 250$), R: replacement. $t = 10^6$, $n = 2$.

This is by no means a formal proof of the catastrophe principle holding for our data, though we do see that the ratios in Table 2 are 1 or close to 1. The long (and plausibly heavy-tailed) time to convergence would be explained by the existence of metastable state in which the process is ‘stuck’ in for a long time before at some point the process jumps to a true steady state in consensus.

⁷It has been shown [8, 11] that if (19) holds for some $n \geq 2$ then it holds for all $n \geq 2$.

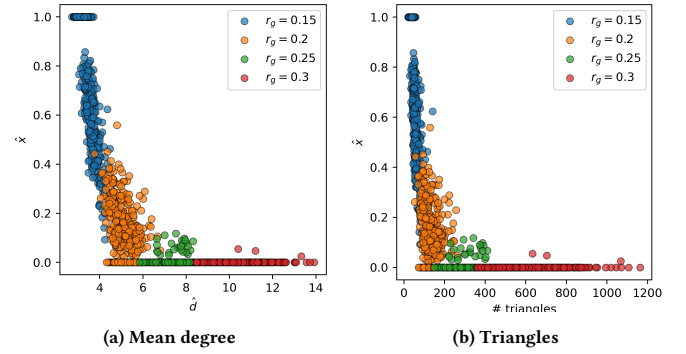


Figure 3: Average final estimate scatter plotted against various network characteristics. Each point is not fully opaque so that darker regions are indicative of more data points.

In Figure 3 we show scatter plots contrasting network characteristics with the final average agent belief. The color represents the setting of r_g . By these plots and Table 1 we see that for $r_g = 0.15$ convergence is either to the always-contribute steady state or does not occur. Conversely for all other settings $r_g \in \{0.2, 0.25, 0.3\}$ consensus was reached only on the always-defect steady state or not at all in the simulated time. This seems like a rather abrupt transition based solely on the tabulated results, however, in the scatter plots we this is somewhat gradual. In general as the r_g increases the final average belief decreases.

In Figure 3a and 3b we see that the mean degree of the network and the number of triangles in the graph have a strong negative correlation to the average agent belief at the end of the simulation run. As the average degree of the network increases, so does the average of the interaction size k_t . A similar relationship holds with the number of triangles and the interaction size. It is not surprising that a bigger interaction size k_t should result in more convergence to the always-defect steady state. This is because as k_t increases, $F(x^{k_t-1})$ decreases and so the per turn probability of contributing decreases.

We depict the network and the belief of agents for illustrative simulation runs that did not converge in simulated time in Figure 4. Networks that do not reach convergence in simulated time have regions which are dense, as well as sparse. In the dense regions, belief is low, and in the sparse regions belief is high. In Figure 4 we see how in the densely connected regions of the network trust is low, while in the sparsely connected regions trust is high. This illustrates that metastable disagreement is a phenomena that emerges with network structure⁸. We know that if we could observe the system for an infinite time, then eventually by chance the beliefs and actions of all agents will converge to either all contribute or all-defect state. In the non-equilibrium states, agent belief is highly dependent on the local network structure.

⁸By the nature of λ taking values in \mathbb{R} , the probability that $x^{k-1} < 1/\lambda$ is increasing in $x \in [0, 1]$ and decreasing in $k \in \{2, N\}$. Thus, regardless of distribution of λ , we expect a cutoff in the group-size under (above) which contribution (defection) is favored. This would result in similar patterns observed in Figures 3, and 4 simply shifted depending on where the cutoff is.

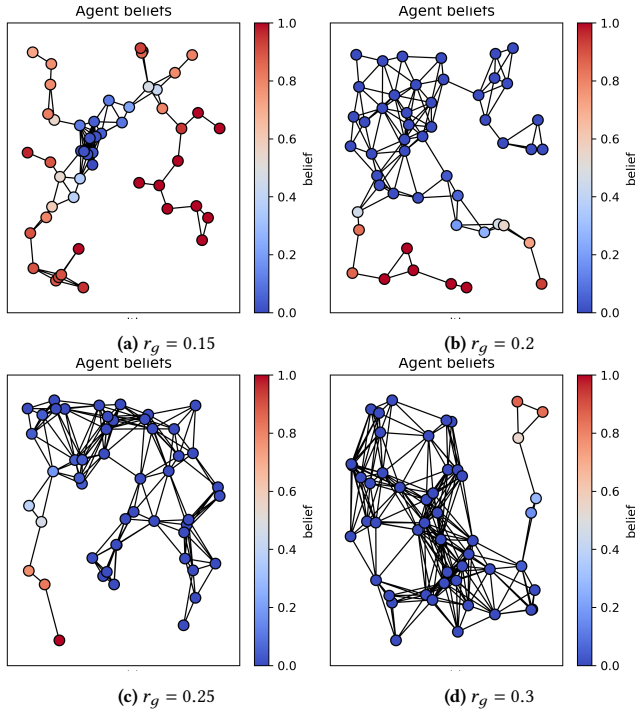


Figure 4: State at termination of simulation for iterations in which convergence was not reached. This illustrates how the agent belief before absorption in a state of complete or zero trust appears localized, with high trust in regions with few connections, and low trust in regions with many connections.

3.4 Regular graphs, and metastability in geometric graphs

To disentangle the effect of the complex network topology and the effect of interaction size, we also study the dynamics on circulant graphs. We consider 50 agents spaced on a circle, and connect each agent to its l nearest neighbors⁹ for $l \in \{2, 4, 6, 8\}$. Each agent has the same degree, allowing us to study the effect of interaction size without heterogeneity in density. We keep the CDF $F(x) = 1/x^4$, and $\alpha = 0.3$. The proportion of simulation runs which converge to the always-contribute and the always-defect steady state is tabulated in Table 3. In contrast to the simulation runs on the random geometric graphs, all the simulation runs converge within the simulated time.

l	2	4	6	8
Contribution	1	0.508	0	0
Defection	0	0.492	1	1

Table 3: Portion of runs out of 500 which converged to contribution and defection in 10^7 rounds.

For the circulant connecting each agent to l nearest neighbors, the games played by the agents are always of size $k = l + 1$. Thus for $l = 4$ the an agent with belief x contributes with probability x . Conversely when $l < 4$ ($l > 4$) an agent with belief x cooperates at probability $y > x$ ($y < x$) thus favoring contribution (defection) which is reflected in the results of the simulation tabulated in Table 3.

⁹These correspond to the circulants C_{50}^1 , $C_{50}^{1,2}$, $C_{50}^{1,2,3}$ and $C_{50}^{1,2,3,4}$.

The tail probabilities of the time to convergence τ for the regular graphs are plotted in Figure 5 on a log-log scale. In this case τ is not heavy-tailed. Furthermore, as l increases from 2 to 4 the time to convergence increases, and then decreases again from 4 to 6 and 6 to 8. We expect that for $F(x) = 1/x^m$ a similar pattern would emerge with an increase until $l = m$, and a subsequent decrease. The dynamics are fastest when the games are either smallest or largest. At the threshold $l = m$ they would be slowest to converge, because then neither steady-state is favored, as illustrated in Table 3.

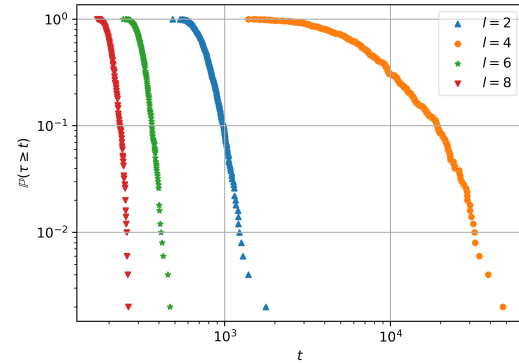


Figure 5: Tail probabilities ($\mathbb{P}(\tau \geq t)$) (on a log-log scale) for the time to convergence for different values of nearest neighbors l .

The difference in tail distributions seen when comparing Figure 2 and Figure 5 indicates the behavior of the dynamics is fundamentally different in these two cases. We posit this difference is due to metastable states when the population structure takes the random geometric graph form. To test and illustrate this we run additional simulation runs paying attention to the total belief of the agents in the population. If a state of the agent beliefs is metastable we expect $\sum_{i=1}^N x_i(t)$ to reach some value and subsequently fluctuate (due to the randomness in the system) around this value before eventually being absorbed in one of the two states $\sum_{i=1}^N x_i(\tau) \approx 0$ or $\sum_{i=1}^N x_i(\tau) \approx N$. When running these simulations, we change the maximum running time to 10^6 and vary $\alpha \in \{0.1, 0.3\}$. We choose a shorter running time to avoid overly long simulation runs. We vary the value of α to illustrate that the behavior is qualitatively the same.

We plot the total population belief over time for two runs (one each for $\alpha = 0.1$ and $\alpha = 0.3$) which illustrate metastable behavior in Figure 6. The total belief quickly reaches a value around which fluctuation occur for a finite time before absorbing in the steady state where $x \approx 1^N$. Observe that the fluctuations are larger for bigger α . This is sensible because α is the learning step size; how much the random fluctuations of the most recent round weigh in the agent belief.

We plot the network and agent beliefs of a representative round during the metastable period of these two illustrative runs in Figure 7. These states are similar to the states plotted in Figure 4 in having a region with low density (smaller game size and more trust) and a region with high density (larger game size and less trust). This is not possible in regular graphs and thus the most likely explanation for the metastability observed in random geometric networks.

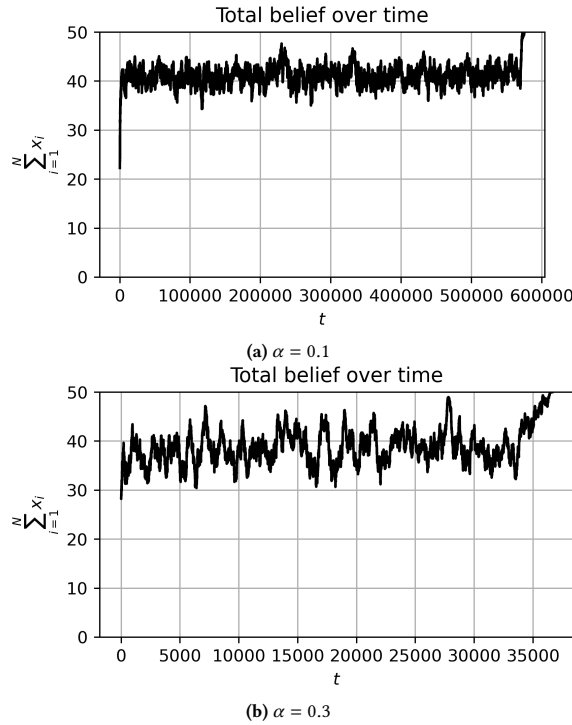


Figure 6: The total agent belief over time for simulation runs in which we posit a metastable state is reached in finite time and eventually exited before $t = 10^6$. In these simulations $N = 50$, $r_g = 0.15$ and α is indicated in the subcaption.

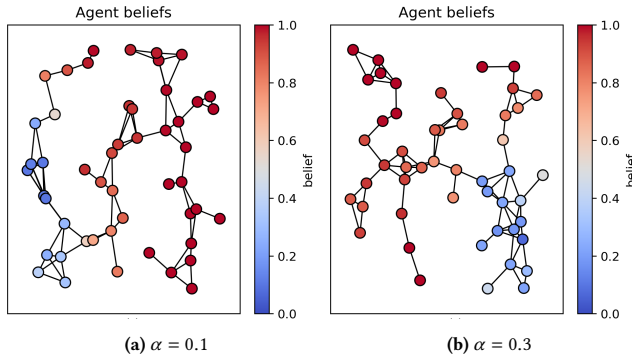


Figure 7: Representative agent beliefs from the metastable period illustrated in Figure 6. In these simulations $N = 50$, $r_g = 0.15$ and α is indicated in the subcaption.

4 CONCLUSION

4.1 Summary and discussion

In this paper we have defined a model for an all-or-nothing version of the public goods game with random payoffs to be played on networks. The agents in the model use the adaptive learning rule called exponential moving average to estimate the behavior of the agents in their neighborhood of the network. This learning rule is adaptive because the weight given to the most recent observation does not diminish with time.

The main theoretical result of the paper is Proposition 3.1 which states that eventually a steady state must be reached. The steady state in question either involves all agents always playing the contribute action or all agents always playing the defect action.

By simulation we study the effect of complex networks on the time and destination of this convergence. In particular we studied the effect of the radius in the random geometric graph model. We see that a larger radius (which implies more connections) results in more convergence to the defect steady state. In agreement with this, when simulating the dynamics on regular graphs, a larger interaction size results in more runs resulting in long-term defection. In comparison to the regular network, the complex network structure slows down convergence, as recently demonstrated for opinion dynamics [38].

By the nature of the distribution of the time to convergence we posit that the system admits metastable states in which both contribution and defection take place. Iterations which do not converge in the simulated time have heterogeneous topology including regions of high and low density. This is confirmed when we inspect the total (summed) agent belief which may reach metastable value (with fluctuations) for a long period of time before ‘jumping’ to one of the truly stable states.


4.2 Future work

The theoretical result in this paper is only proved for the exponential moving average learning rule. It would be good contribution to classify learning rules that do and do not result in asymptotic convergence to steady state of a pure strategy on connected networks. We suspect that when the underlying game admits pure Nash equilibria, these may be learned by a variety of algorithms also in the network setting. Conversely it is interesting to investigate whether changing the learning rate α to be decreasing with t would stabilize the states posited to be metastable in this paper.

We posit metastability of the non-converged dynamics in this model and have illustrated this with examples. With more computational power it may be possible to conduct experiments which elucidate the reasons behind the possibility of metastable non-convergence. The question arises: ‘What network characteristics (local or global) are required to allow metastable states with positive probability?’ We suggest looking into things such as heterogeneity of the local characteristics in the graph.

A massive body of literature exists on evolutionary dynamics on regular graphs such as grids, and tori. In this paper we observe very different behavior on geometric networks (intermediate disagreement) and regular networks (fast convergence to steady-state). This highlights how important it is to consider the effects of non-regular network topology on learning and evolutionary dynamics. The networks considered in this paper are highly stylized, thus illustrating a proof of concept. It is left for future work to delve deeper into the consequences of more realistic network structures.

ACKNOWLEDGMENTS

This research was supported by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement no. 945045, and by the NWO Gravitation project NETWORKS under grant no. 024.002.003. 

REFERENCES

- [1] S Banisch and E Olbrich. 2019. Opinion polarization by learning from social feedback. *The Journal of Mathematical Sociology* 43, 2 (2019), 76–103. <https://doi.org/10.1080/0022250X.2018.1517761>
- [2] H Barcelo and V Capraro. 2015. Group size effect on cooperation in one-shot social dilemmas. *Scientific Reports* 5, 1 (2015), 7937. <https://doi.org/10.1038/srep07937>
- [3] W Barfuss, J F Donges, V V Vasconcelos, J Kurths, and S A Levin. 2020. Caring for the future can turn tragedy into comedy for long-term collective action under risk of collapse. *Proceedings of the National Academy of Sciences* 117, 23 (2020), 12915–12922. <https://doi.org/10.1073/pnas.1916545117>
- [4] F Bierbrauer and J Winkelmann. 2020. All or nothing: State capacity and optimal public goods provision. *Journal of Economic Theory* 185 (2020), 104955. <https://doi.org/10.1016/j.jet.2019.104955>
- [5] A J Bladon and T Galla. 2011. Learning dynamics in public goods games. *Physical Review E* 84 (2011), 041132. <https://doi.org/10.1103/PhysRevE.84.041132>
- [6] T Börgers and R Sarin. 1997. Learning through reinforcement and replicator dynamics. *Journal of Economic Theory* 77, 1 (1997), 1–14. <https://doi.org/10.1006/jeth.1997.2319>
- [7] X-B Cao, W-B Du, and Z-H Rong. 2010. The evolutionary public goods game on scale-free networks with heterogeneous investment. *Physica A: Statistical Mechanics and its Applications* 389, 6 (2010), 1273–1280. <https://doi.org/10.1016/j.physa.2009.11.044>
- [8] V Chistyakov. 1964. A theorem on sums of independent positive random variables and its applications to branching random processes. *Theory of Probability and its Applications* 9 (1964), 640–648. <https://doi.org/10.1137/1109088>
- [9] J Dall and M Christensen. 2002. Random geometric graphs. *Physical Review E* 66 (2002), 016121. Issue 1. <https://doi.org/10.1103/PhysRevE.66.016121>
- [10] A Dannenberg, A Löschel, G Paolacci, C Reif, and A Tavoni. 2015. On the provision of public goods with probabilistic and ambiguous thresholds. *Environmental and Resource Economics* 61, 3 (2015), 365–383. <https://doi.org/10.1007/s10640-014-9796-6>
- [11] P Embrechts, C Klüppelberg, and T Mikosch. 1997. *Modelling extremal events: For insurance and finance*. Vol. 33. Springer Verlag, Berlin, Germany. <https://doi.org/10.1007/978-3-642-33483-2>
- [12] J Gómez-Gardeñes, D Vilone, and A Sánchez. 2011. Disentangling social and group heterogeneities: Public Goods games on complex networks. *Europhysics Letters* 95, 6 (2011), 68003. <https://doi.org/10.1209/0295-5075/95/68003>
- [13] R M Isaac, J M Walker, and A W Williams. 1994. Group size and the voluntary provision of public goods: Experimental evidence utilizing large groups. *Journal of Public Economics* 54, 1 (1994), 1–36. [https://doi.org/10.1016/0047-2727\(94\)90068-X](https://doi.org/10.1016/0047-2727(94)90068-X)
- [14] M A Javarone and F Battiston. 2016. The role of noise in the spatial public goods game. *Journal of Statistical Mechanics: Theory and Experiment* 2016, 7 (2016), 073404. <https://doi.org/10.1088/1742-5468/2016/07/073404>
- [15] Y M Kaniovski and H P Young. 1995. Learning dynamics in games with stochastic perturbations. *Games and Economic Behaviour* 11 (1995), 330–363. <https://doi.org/10.1006/game.1995.1054>
- [16] R Lv, J-L Qian, Q-Y Hao, C-W Wu, N Guo, and X Ling. 2024. The impact of reputation-based heterogeneous evaluation and learning on cooperation in spatial public goods game. *Chaos, Solitons and Fractals* 181 (2024), 114668. <https://doi.org/10.1016/j.chaos.2024.114668>
- [17] G J Mailath and A Postlewaite. 1990. Asymmetric information bargaining problems with many agents. *The Review of Economic Studies* 57, 3 (1990), 351–367. <https://doi.org/10.2307/2298018>
- [18] M McGuire. 1974. Group size, group homogeneity, and the aggregate provision of a pure public good under Cournot behavior. *Public Choice* 18 (1974), 107–126. <https://doi.org/stable/30022752>
- [19] M McPherson, L Smith-Lovin, and J M Cook. 2001. Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology* 27 (2001), 415–444. <https://doi.org/10.1146/annurev.soc.27.1.415>
- [20] B V Meylahn, A V den Boer, and M R H Mandjes. 2024. Interpersonal trust: Asymptotic analysis of a stochastic coordination game with multi-agent learning. *Chaos* 34, 6 (2024), 063119. <https://doi.org/10.1063/5.0205136>
- [21] B V Meylahn and J M Meylahn. 2024. How social reinforcement learning can lead to metastable polarisation and the voter model. *PLoS ONE* 19, 12 (2024), e0313951. <https://doi.org/10.1371/journal.pone.0313951>
- [22] J Nair, A Wierman, and B Zwart. 2022. *The Fundamentals of Heavy Tails: Properties, Emergence, and Estimation*. Cambridge University Press, Cambridge, UK. <https://doi.org/10.1017/9781009053730>
- [23] D Nosenzo, Q Simone, and M Sefton. 2015. Cooperation in small groups: The effect of group size. *Experimental Economics* 18, 1 (2015), 4–14. <https://doi.org/10.1007/s10683-013-9382-8>
- [24] M Penrose. 2003. *Random Geometric Graphs*. Oxford University Press, Oxford, UK. <https://doi.org/10.1093/acprof:oso/9780198506263.001.0001>
- [25] M Perc. 2007. Transition from Gaussian to Levy distributions of stochastic payoff variations in the spatial prisoner's dilemma game. *Physical Review E* 75 (2007), 022101. Issue 2. <https://doi.org/10.1103/PhysRevE.75.022101>
- [26] J-L Qian, Y-X Zhou, and Q-Y Hao. 2024. The emergence of cooperative behavior based on random payoff and heterogeneity of concerning social image. *Chaos, Solitons & Fractals* 179 (2024), 114405. <https://doi.org/10.1016/j.chaos.2023.114405>
- [27] F C Santos, M D Santos, and J M Pacheco. 2008. Social diversity promotes the emergence of cooperation in public goods games. *Nature* 454 (2008), 213–216. <https://doi.org/https://doi.org/10.1038/nature06940>
- [28] Y Sato, E Akiyama, and J P Crutchfield. 2005. Stability and diversity in collective adaptation. *Physica D: Nonlinear Phenomena* 210, 1 (2005), 21–57. <https://doi.org/10.1016/j.physd.2005.06.031>
- [29] Y Sato, E Akiyama, and J D Farmer. 2002. Chaos in learning a simple two-person game. *Proceedings of the National Academy of Sciences* 99, 7 (2002), 4748–4751. <https://doi.org/10.1073/pnas.032086299>
- [30] Y Sato and J P Crutchfield. 2003. Coupled replicator equations for the dynamics of learning in multiagent systems. *Physical Review E* 67 (2003), 015206. Issue 1. <https://doi.org/10.1103/PhysRevE.67.015206>
- [31] S Suzuki and E Akiyama. 2005. Reputation and the evolution of cooperation in sizable groups. *Proceedings of the Royal Society B* 272 (2005), 1373–1377. <https://doi.org/10.1098/rspb.2005.3072>
- [32] A Szolnoki and M Perc. 2010. Impact of critical mass on the evolution of cooperation in spatial public goods games. *Physical Review E* 81, 5 (2010), 057101. <https://doi.org/10.1103/PhysRevE.81.057101>
- [33] A Szolnoki and M Perc. 2011. Group-size effects on the evolution of cooperation in the spatial public goods game. *Physical Review E* 84, 4 (2011), 047102. <https://doi.org/10.1103/PhysRevE.84.047102>
- [34] A Tavoni, A Dannenberg, G Kallis, and A Löschel. 2011. Inequality, communication, and the avoidance of disastrous climate change in a public goods game. *Proceedings of the National Academy of Sciences* 108, 29 (2011), 1185–11829. <https://doi.org/10.1073/pnas.1102493108>
- [35] J Wang, F Fu, T Wu, and L Wang. 2009. Emergence of social cooperation in threshold public goods games with collective risk. *Physical Review E* 80 (2009), 016101. <https://doi.org/10.1103/PhysRevE.80.016101>
- [36] X Wang, Z Yang, G Chen, and Y Liu. 2024. Enhancing cooperative evolution in spatial public goods game by particle swarm optimization based on exploration and q-learning. *Appl. Math. Comput.* 469 (2024), 128534. <https://doi.org/10.1016/j.amc.2024.128534>
- [37] J Weimann, J Brosig-Koch, T Heinrich, H Hennig-Schmidt, and C Keser. 2022. CO2 Emission reduction – Real public good provision by large groups in the laboratory. *Journal of Economic Behavior & Organization* 200 (2022), 1076–1089. <https://doi.org/10.1016/j.jebo.2022.06.027>
- [38] F Zarei, Y Gandica, and L E C Rocha. 2024. Bursts of communication increase opinion diversity in the temporal Deffuant model. *Scientific Reports* 14 (2024), 2222. <https://doi.org/10.1038/s41598-024-52458-w>
- [39] Z Zeng, Q Li, and M Feng. 2022. Spatial evolution of cooperation with variable payoffs. *Chaos* 32, 7 (2022), 073118. <https://doi.org/10.1063/5.0099444>