

Co-evolution of in-group favoritism and cooperation

Diogo Abreu

diogo.perestrelo.abreu@tecnico.ulisboa.pt

Instituto Superior Técnico, Lisboa, Portugal

October 2018

Abstract

In-group favoritism is the tendency of individuals to aggregate towards others they have more in common with and adapt their behavior accordingly. It is a central aspect of the real world that strongly impacts the propensity of individuals to cooperate, both in human and animal societies: Individuals are often more willing to incur a cost to convey a benefit to individuals of the same group. Here we show that this outcome may depend on the particular dilemma individuals face and how individuals revise their group membership, as well as their attitude towards others, who can share or not this membership. We develop a model that simulates a population of agents that are divided in groups and interact with each other, playing 2-player games which are modeled according to the principles of Game Theory. Our purpose was to study the behavior of the population and how cooperation depends on characteristics such as number of groups, tendency towards in-group or out-group interaction (which can be expressed as the proportion of agents who tend to cooperate more with outsiders), as well as the game that is played. We studied their impact on strategies, group membership, and overall cooperation, be it within each group the overall population. We show that whenever group affiliation remains unchanged, irrespectively of the game played, increasing the number of groups raises cooperation. Group imitation promotes cooperation in coordination games such as the Stag Hunt game, but has the opposite effect with the Prisoners Dilemma. Finally, by increasing the proportion of agents that interact less with outsiders, we obtain higher values for overall cooperation in the Prisoners Dilemma, whereas the Stag Hunt sees higher values of cooperation when either all agents interact with outsiders, or none at all.

Keywords: In-Group Favoritism, Game Theory, Cooperation, Multi-Agent System, Learning

1. Introduction

The concept of groups is fundamental in the interaction between human beings, serving as a basis for all cultures and extending to other living organisms. These groups are defined by an abundance of interactions between members of the same group and a difference of behavior when dealing with these versus members of other groups. This bias towards in-group interaction can be referred to as In-group Favoritism or Parochialism, and in order to understand it, it is useful to frame it as a problem of cooperation.

Cooperation can be defined[19] as a situation in which an agent performs an action which has a cost c , so another agent may receive a benefit b , where $b > c$.

In a population made up of agents which cooperate and others that defect, we would expect that the cooperators would be at a disadvantage, given that they pay a cost for their actions that defectors do not. Defectors over time should take advantage of the cooperators and ultimately become more prevalent in a population. This, however, is not what

we observe in the real world, where cooperation is abundant. It is therefore the problem of Cooperation to understand how it emerges given that there is no apparent reason for an agent to choose cooperation over defection. Here we will study this phenomenon and try to answer the following questions:

- How do agents act when playing different games under conditions such as group membership, strategy and group imitation, and social learning?
- What role does group membership play in the evolution of cooperation and in-group favoritism?
- How is this evolution affected by the agents' willingness to interact outside their group?
- Does changing the proportion of agents willing to interact with outsiders have an effect on cooperation? And does it produce different results in different games?

The last decades have witnessed the discovery of several mechanisms that allow for the emergence of cooperation in a population of self-regarding individuals: direct and indirect reciprocity, multi-level and kin selection, networked population, signaling, punishment, among many others. Direct reciprocity [35, 10, 11, 36, 17] is a mechanism that appears as a result of repeated interactions between two individuals, where the probability of another interaction is sufficiently high. It allows for the emergence of cooperation when we consider that it may be worth paying the cost now to warrant cooperation by the other in the future.

Indirect reciprocity [16, 14, 13, 27] introduces the concept of reputation. The interaction between two individuals is observed by others, and the information regarding the outcome spreads throughout the population. Individuals can therefore base their actions on the reputation of those with whom they are interacting, as well as with the interest of increasing their own reputation. The cost of cooperation in this case may be low in comparison with the cost of a reduced reputation, which implies less benefits from others in the future. This mechanism contributes to the emergence of cooperation if the probability of knowing someone's reputation is sufficiently high.

When describing a population based on the probability of interaction between its individuals, we can say that the population is well-mixed, in which case every agent is equally likely to interact with every other agent, or there can be a structure to a population, in which there is a bias towards interaction with certain individuals. If we consider a spatial structure in which agents are more likely to interact with their neighbors, then the evolution of cooperation may be dependent on the appearance of clusters of cooperators [12, 15, 20], whose members receive higher benefits than those who share a cluster with one or more defectors. This effect seems to be amplified [23, 31, 25] in populations which interact through heterogeneous networks, which are ubiquitous in social and biological systems [2].

These interactions can exist not only between individuals, but also between groups. Multi-level selection [37, 3, 33] describes competition between groups of individuals. This mechanism further promotes the idea that the success of cooperators comes from groups, since a group of cooperators will have a bigger total fitness than a group with defectors.

Instead of looking at groups or individuals, we can also look at characteristics, or genes. Kin selection [6] is a mechanism that operates if an individual's actions are conditioned to the recognition of his kin. The primary motivation is the survival of characteristics common among individuals. The

cooperation between individuals sharing characteristics is simply one or more genes maximizing their fitness and ensuring their survival.

The five previous mechanisms when combined make a strong case for the possibility of the emergence of cooperation. Other interaction structures, however, were also pointed as central in promoting and sustaining cooperation.

The role of punishment has been studied [28, 4, 8] in regards to its capacity to promote cooperation within a population when public goods games are played. In these games, members of a group engage in a common activity where a public good is shared by all. Cooperators contribute to the public good by paying a cost, and defectors reap the benefits of this good with no contribution. By allowing individuals to punish those who contribute less, cooperation can emerge.

The use of signals is another way in which cooperation has been shown to emerge [1, 30, 24].

Signaling involves a sender, a receiver and a signal that is sent in between. The sender analyzes the state of the environment which they inhabit, sending a signal to the receiver who as a result performs an action from a predetermined set. Each action has an associated payoff. There is no inherent meaning to the signals, they are sent arbitrarily to the receiver. It is interesting to see that in these chaotic environments signals can gain meaning and foster cooperation. A signaling network is, therefore, another mechanism that can promote cooperation.

1.1. In-group Favoritism

As with the cases previously described, we shall study how in-group favoritism can evolve alongside cooperation and how it can promote its emergence within a population. A proper description of the interaction between individuals can be made applying concepts from Game Theory, which we will now present.

The interaction between agents will be modeled according to principles of Game Theory. Agents will play one-shot games against each other. In a simulation we choose two paradigmatic 2-person symmetric games to be played: the prisoner's dilemma [18], a defection dominance dilemma, and the Stag Hunt game [29], which describes a coordination problem. These two games and their respective payoffs can be modeled with a simple payoff matrix, represented in Fig. 1.

Each entry of the payoff matrix represents a payoff for the player whose plays are represented by the matrices' lines, in this case player A, according to the different combination of both players' choices in the game they're playing. R represents mutual cooperation and P mutual defection. If $P = 0$ and $R = 1$, then we can adjust the relative values of T

		Player B	
		C	D
Player A	C	R	S
	D	T	P

Figure 1: Payoff Matrix

(Temptation to Defect) and S (Sucker’s Payoff) in order to represent one of the two games.

Provided that mutual cooperation is always preferred over mutual defection, the two dilemmas arise naturally, depending on the relative ordering of these four payoffs: the Stag-Hunt game (SH), for which $1 > T > 0 > S$; and the Prisoners Dilemma game (PD), for which $T > 1 > 0 > S$.

For both dilemmas, mutual cooperation advantageous over unilateral cooperation and mutual defection. Tension will arise when players prefer mutual defection to unilateral cooperation ($S < 0$). Here, individuals tend to defect because they fear that their opponent will also defect. For this reason, the difference $|S| - P$ is often coined as "fear". If besides that, players prefer unilateral defection to mutual cooperation ($T > R = 1$), individuals have an additional incentive to defect, as they try to cheat their opponent. The difference $T - R$ is often coined as greed.

Following this terminology, the SH is based on fear; the PD combines both fear and greed. In the SH individuals try to coordinate actions: cooperate when the other cooperates, and defect if the opponent defects. In the PD, the rational choice would be to defect, irrespectively of the other’s choice, being therefore often referred to as a defection dominance dilemma.

Most of the literature focuses on the PD as the most paradigmatic game of cooperation. More recently, several authors have argued that the SH represents a more significant context in which to study cooperation [29]. In fact, in many circumstances, it is often hard to differentiate the two, as it depends on how the fitness is calculated. Here we will address both the PD and SH, highlighting the differences among the two.

Taking these concepts into consideration, and building upon the work presented in the next section, we pretend to develop a model where a population of agents is divided into groups and play games one-on-one, in which they can cooperate or defect against their opponent. These games are modeled according to principles from Game Theory. Agents will have strategies which tell them if they should cooperate or defect, can adopt other agents’ groups and strategies, and have a bias in the proportion of interactions with in-group and out-group members.

This bias can be altered to allow agents to interact more with outsiders than with those of their own group. This bias can also be different for different agents, and we will explore how this difference affects cooperation and favoritism in the population. This paper is organized as follows: in Section 2 we talk about the related work that serves as a basis for our own; Section 3 is dedicated to the description of the architecture and implementation of our system; Section 4 presents our results; In Section 5 we draw our conclusions.

2. Related Work

In order to implement a system that would allow us to explore the concepts of in-group favoritism and cooperation within a population, a better understanding of all aspects of the system was needed. This understanding was acquired through related work which presents other systems which explore concepts such as the evolution of in-group favoritism, the effects of separating a population into groups, using tags as a way to create these groups and how agents can learn to adapt their behavior within this environment.

Of all the works studied, the one that provided the basis for our work and most the concepts we used is titled **The evolution of in-group favoritism**[5]. In this paper a mathematical framework for the analysis of the evolution of in-group favoritism within a population is created. Agents within a population interact with each other by playing the Prisoner’s Dilemma. Since agents are divided into groups, a different behavior can be adopted when interacting with members of the same group versus those of others. This behavior is defined by a strategy made up of two values $[p, q]$, where p is the probability of cooperating with a member of the same group, and q the probability of cooperating with members of other groups. If $p > q$ then cooperation is in-group biased; for $p < q$ cooperation is said to be out-group biased. Agents can not only adopt each other’s strategies but also each other’s groups. The results obtained in the article are the outcome of individuals playing a Prisoner’s Dilemma game against each other, with each simulation presenting different values for number of groups M , cost of cooperation c , strategy mutation rate μ and migration rate between groups v . Another important metric used to understand the emergence of in-group favoritism is the ratio between the benefit and cost for cooperation, $\frac{b}{c}$.

To better understand and model in and out group interactions, we referred to the paper titled **The evolutionary interplay of intergroup conflict and altruism in humans: a review of parochial altruism theory and prospects for its extension**[22]. This paper provides a solid definition of these interactions, as well as defining

the behavior of individual agents in relation to its effects on the one performing the action and the recipient. From here we can clearly define in-group favoritism, and the concept of parochial altruism, the idea that the readiness to behave altruistically to the benefit of in-group members and to act hostilely towards out-groups are two closely linked aspects in human evolution.

Another important aspect to understand is how to create groups, and how they have behaved in other works. A. Traulsen [32] in his work introduces the concept of tag-based cooperation, where individuals cooperate based on an arbitrary tag. Here two models for tag cooperation are considered. The first model, based on work by Jansen and van Baalen[9], is termed *cooperative tags*. It considers the creation of groups of cooperators based on a common signal. The term *cooperative tags* is used due to the fact that tags are used as a way to establish cooperation in a sea of defectors.

The second model, based on work by Riolo et al.[21], focuses on the creation of self-serving groups of cooperators, where the infiltration of defectors is fatal, therefore it is termed *defective tags*.

Another important work related to tag-based cooperation is by Traulsen and Nowak [34], which presents another approach to the evolution of tag-based cooperation. The analogy used in this case is that of each tag being a color that allows the establishment of new signals of recognition, hence the title **Chromodynamics of Cooperation in Finite Populations**. Here, as in other works, there is a benefit and a cost for cooperation, represented by b and c , respectively. The game played by the individuals in the population is again the Prisoner's Dilemma, and thus they can again be classified as cooperators or defectors. Cooperators may recognize each other by their tag, which can be seen as a secret handshake. However, once this handshake is discovered by a defector, it loses its purpose and a new handshake must be agreed upon. There is a constant race between these two classes: cooperators constantly try to establish new handshakes and defectors try to decipher them. This paper provides a good visualization of group migration as a result of the use of tags. Finally we have a paper titled **The evolution of ethnocentrism**[7], which cements the effects of tags in the development of in-group favoritism, albeit with spatial restrictions in this case.

The final aspect studied in the related work was learning. In an article titled **Dynamics of Fairness in Groups of Autonomous Learning Agents**[26] we learn about the concept of fairness applied to an n -player game. Learning to play this game can be done in one of two ways: social learning or individual learning. The focus of this work is on

individual learning, namely on the implementation of the Roth-Erev reinforcement learning algorithm. This algorithm is used to analyze the outcome of a population of learning agents that play a multi player ultimatum game, by mimicking the learning process of humans while playing social dilemmas. Even though this algorithm is not as relevant to our work, the idea that individual and social learning may yield more appropriate results when simulating populations is of great importance to us.

Having discussed the basis on which our work is built, we now proceed to the description of our own model on the next section.

3. Implementation

Our system was implemented using the C++ language, running on a Ubuntu environment. The UML in Figure 2 describes the architecture of the system.

We consider a population of N agents, each belonging to one of M groups, who will interact through a game in a stochastic environment. The game is a pairwise interaction in which both participants can make one of two choices: cooperate or defect. Each agent has a strategy, described by two values, one for the probability of cooperating with members of the same group, p , and the other for the probability of interacting with members of other groups, q . At the beginning of every simulation, we attribute a random value for p and q for every agent. These values will change as the agents learn and as some strategies seem more advantageous than others.

We consider the population to be well-mixed, however there is a bias which can be set to favor a higher frequency of interaction with members of the same group or with those of other groups. The probability of interacting with members of the same group is referred to as α_{in} , and the probability of interacting with members of other groups, α_{out} . This probability α_{in} can differ from agent to agent, being high in some and low in others, or be the same for all agents.

The fitness function of an agent is the sum of the payoffs of all the games played in one simulation. So the fitness function of an agent i after m interactions with a payoff p_{ij} is given by the following expression:

$$f_i = \sum_j^m p_{ij}$$

In order to allow agents to adapt their strategy, we develop a mechanism of social learning, where agents can imitate each other's strategies with an associated error with a probability obtained from the following expression:

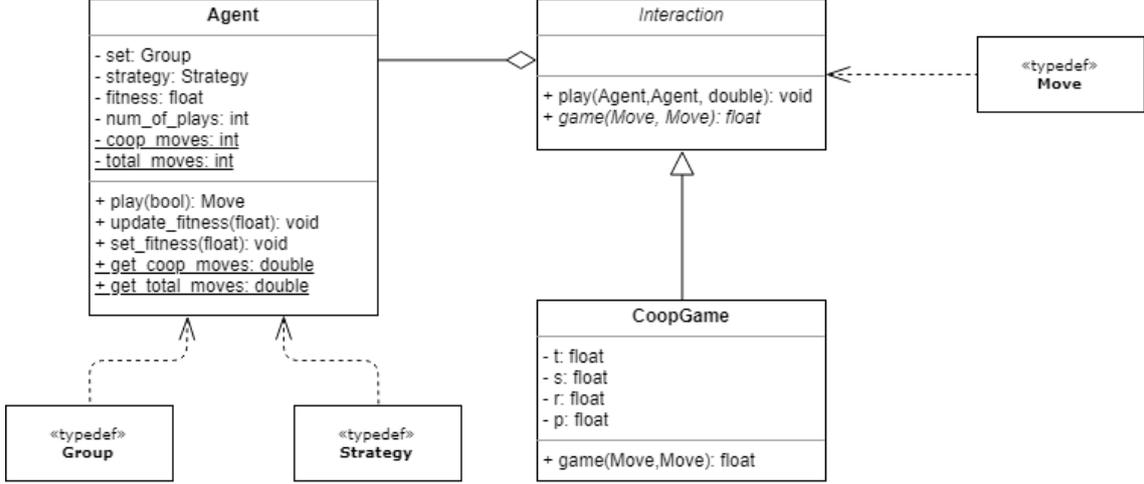


Figure 2: UML for our system. The abstract class *Interaction* allows for future development of the system by allowing the implementation of other games. *CoopGame* is a class that implements our version of the games played. The types *Group*, *Strategy*, and *Move* were also created to allow for future extensions.

$$c_{XY} = \frac{1}{1 + e^{-\beta(f_Y - f_X)}}$$

Here β defines the intensity of selection, f_x and f_y represent the payoffs of each agent. It is important to note that playing the game and adopting a strategy are two disassociated actions. An agent learns from another by selecting one agent at random and comparing the resulting agents fitness function, f_y , to his own, f_x , and therefore obtaining a value for the probability of adopting the strategy. The adoption of a strategy is not perfect, as a deviation of between 0 and 0.01 exists that is added to the copied values of p and q .

Agents can also imitate each other's groups. This happens with a probability that is equal to the probability of interacting with agents from other groups. The idea is that agents will use a similar criterion when deciding to interact with outsiders as they do when choosing to adopt their group. When adopting a group, an agent also has a probability of adopting a random group instead.

3.1. Simulation

As we can see in the previous algorithm, a simulation consists of creating the agents, attributing random values to p and q for all, and having them interact with each other for a defined number of cycles. In order to run a simulation, the following parameters need to be set:

- Value of S ;
- Value of T ;
- Number of groups;
- Number of agents;

- Proportion of in-group interactions α_{in} ;
- value of intensity of selection β ;
- P and Q deviation;
- Probability of adopting a random group;
- proportion of agents with high α_{in} ;
- Number of cycles.

In order to reduce the number of possibilities for testing purposes, we fixed some of these values. The intensity of selection was fixed at 1. Different values were chosen but didn't add depth to the results, and so this seemed like a decent choice. The same mind set went into the choice for number of agents, which was fixed at 1000. This value provides a diverse enough population so that system properties are visible and run time isn't too long, as more agents would imply more cycles so that all agents would be given enough opportunities to interact. P and Q deviation was fixed at 0.01, and random group adoption probability was set at 0.1.

Algorithm 1 describes the interaction cycle, and in it we get a clearer view of how the number of cycles is determined.

This cycle runs N_STEPS times, however this value alone doesn't determine the number of cycles, as the `getFitness()` method calculates the fitness of the agent by having it play N_TRAIN games against other agents. Since we run the `getFitness()` method twice in the interaction cycle, we get a total number of cycles equal to $2N_TRAIN \times N_STEPS$. All of our simulations were done with N_STEPS and N_TRAIN equal to 10.000, which is long enough for the results to be representative without the simulation time being impractical.

Algorithm 1: Interaction cycle

```
for  $N\_STEPS$  do
  clear agents' fitness;
  a1=randomAgent();
  getFitness(a1);
  a2=randomAgent();
  getFitness(a2);
  p_adopt_strat= $c_{a1a2}$ ();
  strat_rand=randomFloat();
  if strat_rand < p_adopt_strat then
    | a1.strategy = a2.strategy;
  end
  group_mutate_rand = randomFloat();
  if group_mutate_rand < p_mutate_group
  then
    | a1.group = randomGroup();
  else
    group_rand = randomFloat();
    if group_rand < p_adopt_group then
      | a1.group = a2.group;
    end
  end
end
end
```

Having described how the system works, in the next section we analyze the results we obtained.

4. Results

Due to the stochastic nature of our system, every result we obtain is the outcome of an average of 30 simulations.

In these simulations our main focus is mainly on the effect of certain values for our variables on the average values of p , q , in group-favoritism, and, out of all choices made by all agents, what percentage was to cooperate. In-group favoritism was initially thought of as the ratio between p and q , however the range of values this allowed for made the analysis of our results less clear. Therefore we opted for defining in-group favoritism not as the ratio but the difference between p and q , reducing our range to values between -1 (absolute out-group favoritism) and 1 (absolute in-group favoritism).

As for the games played, the Prisoner's Dilemma was defined as having $S = -0.05$ and $T = 1.05$. The Stag Hunt game was set as $S = -0.75$ and $T = 0.75$.

We start with a simulation that we consider the baseline, which is that populations with only one group tend towards full defection. Having this result confirmed is a good sign that our system is working as expected. This is visible in the distribution plot of Fig.3, where the horizontal axis represents p , the vertical axis represents q , and agents are placed in the plot according to their values of

p and q . Brighter colors indicate more agents with those values.

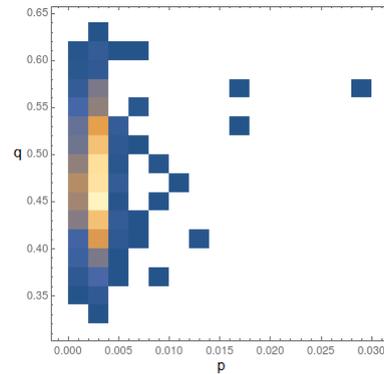


Figure 3: p and q values for 1 group, 1000 agents playing the Prisoner's Dilemma.

The tendency towards full defection is clear, as all agents have a very low value for p . Values of q tend towards 0.5 due to there being no selection for out-group interactions combined with the fact that q values are randomly generated at the beginning of every simulation. There is no use in discussing in-group favoritism when everyone belongs to the same group. Likewise, the values of p tend to 0, which results in no cooperation in the population.

To understand how the number of groups affects this distribution, we kept increasing this number and seeing not only how the average values for p and q were affected, but also the percentage of cooperative moves. This eventually brought us to our first conclusion, that an increase in the number of groups always produces an increase in cooperation, be it in or out group. However, in-group cooperation always seems to out-grow out-group cooperation, so in-group favoritism also tends to increase. These results are the same both for the Prisoner's Dilemma and the Stag Hunt game, as we can see in Fig. 4.

Group imitation was also a parameter that we saw as interesting to study. Previous graphs present us with results that always include group imitation, so our thought was to remove it and see what happens. In figs. 5 and 6 we can see the results.

These results again are interesting due to the different impact the same change has in each of the games. In the case of the Prisoner's Dilemma, removing group imitation was very beneficial for in-group cooperation. However it had the opposite effect on the Stag Hunt game. Thus, our second conclusion: coordination games such as the Stag Hunt game benefit from having group imitation, whereas the Prisoner's Dilemma needs more static groups so that cooperation can increase.

The next parameter we thought of as potentially having a big impact was the proportion of in-group

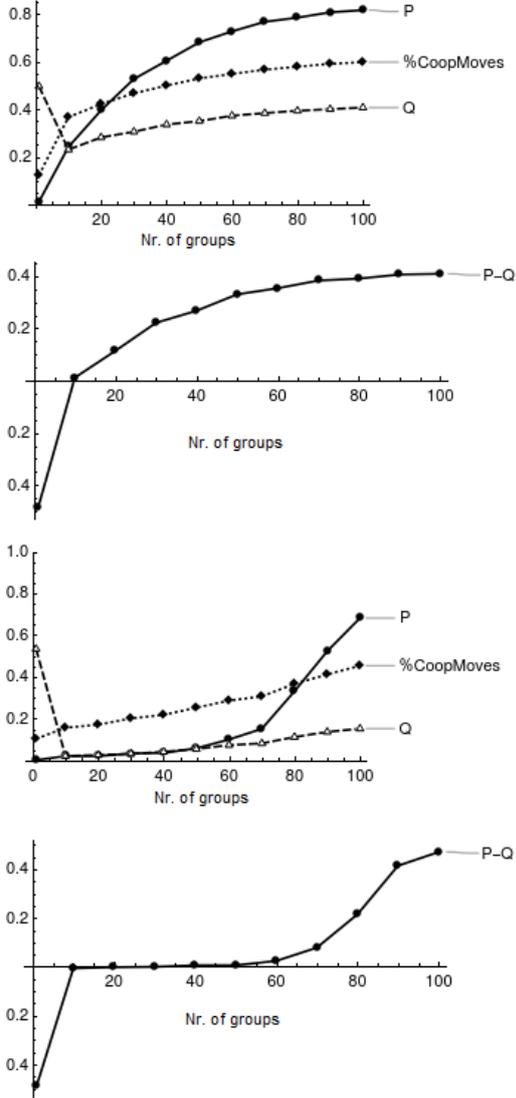


Figure 4: First and third graphs represent p , q , and % of coop moves for between 1 and 100 groups, with 1000 agents playing the Prisoner's Dilemma and Stag Hunt, respectively. The second graph shows the evolution of in-group favoritism for the Prisoner's Dilemma, and the fourth for the Stag Hunt. In all graphs $\alpha_{in} = 0.9$ and group imitation is present.

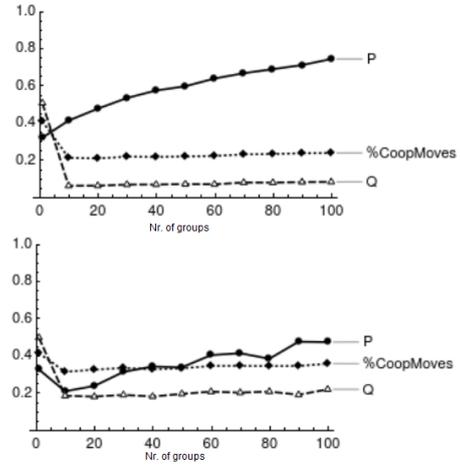


Figure 5: Results for the Prisoner's Dilemma. The top image is a result of removing group imitation. On the bottom we have the original result with imitation. $\alpha_{in} = 0.1$

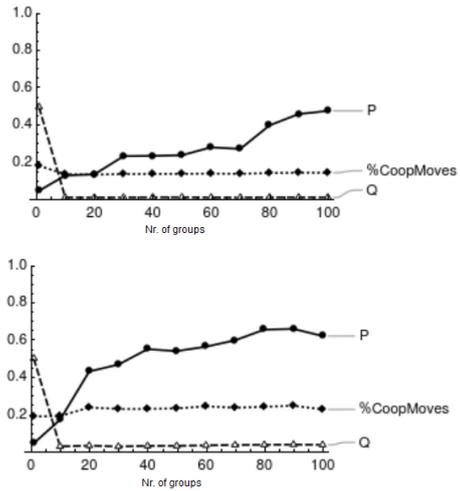


Figure 6: Results for the Stag Hunt. The top image is a result of removing group imitation. On the bottom we have the original result with imitation. $\alpha_{in} = 0.1$

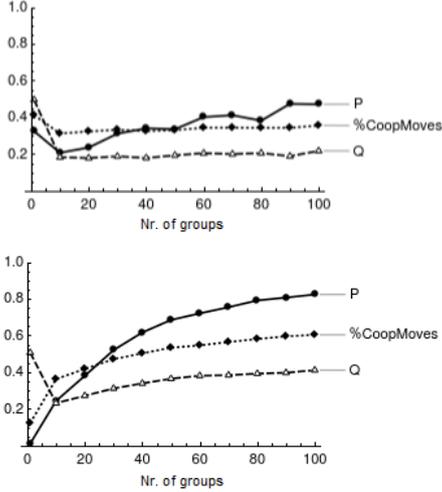


Figure 7: Prisoner’s Dilemma with same parameters as previous graphs. Top represents $\alpha_{in} = 0.1$, and bottom $\alpha_{in} = 0.9$.

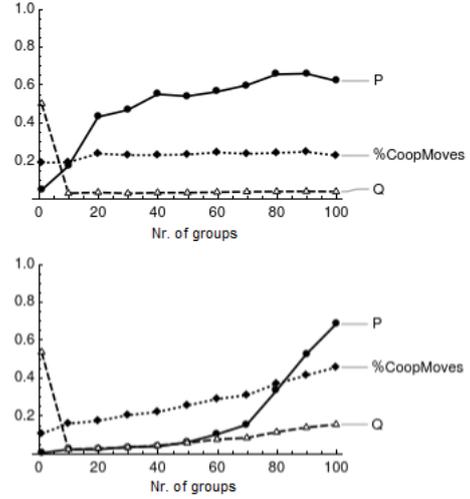


Figure 8: The Stag Hunt game with same parameters as previous graphs. Top represents $\alpha_{in} = 0.1$, and bottom $\alpha_{in} = 0.9$.

interactions, α_{in} . In the previous graphs, this value is set to 0.1. We wanted to see what happened if the reverse was true, so we set α_{in} to 0.9 and compared the results with the previous ones. Figs. 7 and 8 clearly show the impact of this change.

We can see that the lower proportion of in-group interactions has an effect on cooperation. What is interesting is the fact that the Prisoner’s Dilemma sees its cooperation increase significantly, but in the Stag Hunt this tendency to interact more with other groups seems to act as a force against cooperation, only being counteracted after about 70 groups exist. Nevertheless, our hypothesis holds, α_{in} greatly impacts the evolution of cooperation in both games.

There is, however, another aspect to α_{in} that we wanted to explore, the fact that it doesn’t have to be the same for all agents. We can assign a proportion of the population to have a high α_{in} , 0.9, leaving the rest with a α_{in} of 0.1. In this case, we chose as the horizontal axis not the number of groups, which were fixed at 80, but this proportion of agents with high α_{in} . The results are seen in fig. 9.

With these results we draw our final conclusion: a higher proportion of in-group focused agents results in more cooperation for the Prisoner’s Dilemma, but in order to maximize cooperation in the Stag Hunt, we need a homogeneous population in terms of its attitude to interaction with in-group agents. In the next section we summarize our conclusions and discuss future work.

5. Conclusions

The study of cooperation and in-group favoritism is a way of understanding systems that influence many aspects of our lives. If we take games like the Prisoner’s Dilemma or the Stag Hunt game as ways to

model real life interactions, then this understanding can give us more agency over these systems. The following were the main conclusions of our work:

- Increasing the number of groups results in an increase in cooperation.
- Coordination games such as the Stag Hunt benefit from having group imitation, whereas the Prisoner’s Dilemma sees better results without group imitation.
- A higher proportion of agents not willing to cooperate with outsiders promotes cooperation in the Prisoner’s Dilemma. With the Stag Hunt, a more homogeneous attitude is preferred, be it everyone tending to interact with outsiders or no one.

With regards to in-group favoritism, its increase has always followed the increase in the number of groups in all our simulations, and since this is the main mechanism we found that also increases cooperation, this means that promoting cooperation while avoiding this in-group bias seems a quite difficult task.

These conclusions allow us to better understand how agents act when they’re attributed different groups, and how changing aspects such as group imitation and number of groups affects these actions. The willingness to interact outside their group, α_{in} , was also seen to have a great impact on the evolution of in-group favoritism and cooperation. Finally, we were also able to answer the question of whether changing the proportion of agents willing to interact with outsiders would affect these phenomena, and indeed it does.

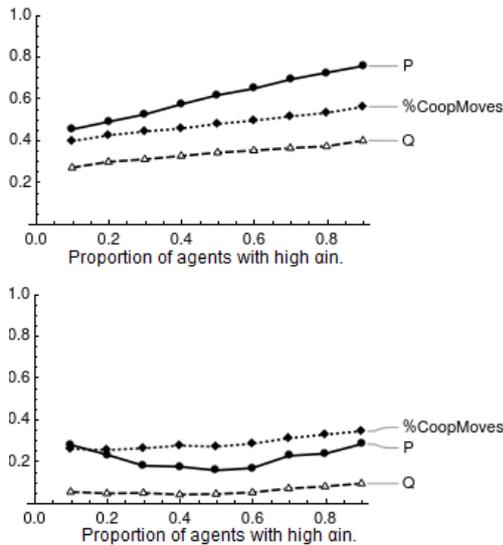


Figure 9: Prisoner’s Dilemma (top) and Stag Hunt game (bottom) with increasing proportion of agents with high α_{in} .

5.1. Future Work

The following points are aspects that we noticed would merit further investigation in future works.

- One of the main points we feel could be better explored in future works is precisely attempting to find the conditions under which cooperation rises without the tribalism associated with in-group favoritism.
- Another aspect is the different effect that α_{in} has on the different games. Adding more depth to this aspect, although desirable on our part, was not given the priority it deserved mainly due to time restraints. It does however seem like a viable path for future studies.
- Since we opted for a mechanism of social learning, a different version of this system could be implemented with individual learning, and the results could be compared with our own.
- The acceptance of members into the group is automatic once the agent decides to join it. This aspect could be improved upon if members of this group could decide to accept or not the new member.

We can see from these suggestions that our work can serve as a basis for a lot of interesting topics relevant to the study of in-group favoritism.

References

[1] R. Argiento, R. Pemantle, B. Skyrms, and S. Volkov. Learning to signal: Analysis of a micro-level reinforcement model. *Stochastic*

Processes and their Applications, 119(2):373–390, 2009.

- [2] A.-L. Barabási. *Network Science*. Cambridge University Press, 2016.
- [3] R. Boyd and P. J. Richerson. Group selection among alternative evolutionarily stable strategies. *Journal of Theoretical Biology*, 145(3):331–342, 1990.
- [4] E. Fehr and S. Gächter. Altruistic punishment in humans. *nature*, 415(6868):137–140, 2002.
- [5] F. Fu, C. E. Tarnita, N. A. Christakis, L. Wang, D. G. Rand, and M. A. Nowak. Evolution of in-group favoritism. *Scientific Reports*, 2012.
- [6] W. Hamilton. The genetical evolution of social behaviour. I. *Journal of Theoretical Biology*, 1964.
- [7] R. A. Hammond and R. Axelrod. The evolution of ethnocentrism. *Journal of Conflict Resolution*, 2006.
- [8] C. Hauert, A. Traulsen, H. Brandt, M. A. Nowak, and K. Sigmund. Via Freedom to Coercion: The Emergence of Costly Punishment. *Science*, 316, 2007.
- [9] V. A. A. Jansen and M. van Baalen. Altruism through beard chromodynamics. *Nature*, 440:663, 3 2006.
- [10] M. Milinski. Tit for tat in sticklebacks and the evolution of cooperation. *nature*, 325(6103):433–435, 1987.
- [11] M. Nowak and K. Sigmund. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner’s Dilemma game. *nature*, 364(6432):56–58, 1993.
- [12] M. A. Nowak and R. M. May. Evolutionary games and spatial chaos. *nature*, 359(6398):826–829, 1992.
- [13] M. A. Nowak and K. Sigmund. Evolution of indirect reciprocity by image scoring. *nature*, 393(6685):573–577, 1998.
- [14] M. A. Nowak and K. Sigmund. Evolution of indirect reciprocity. *nature*, 437(7063):1291–1298, 2005.
- [15] H. Ohtsuki, C. Hauert, E. Lieberman, and M. A. Nowak. A simple rule for the evolution of cooperation on graphs and social networks. *nature*, 441(7092):502–505, 2006.

- [16] H. Ohtsuki and Y. Iwasa. The leading eight: social norms that can maintain cooperation by indirect reciprocity. *Journal of Theoretical Biology*, 239(4):435–444, 2006.
- [17] F. L. Pinheiro, V. V. Vasconcelos, F. C. Santos, and J. M. Pacheco. Evolution of all-or-none strategies in repeated public goods dilemmas. *PLoS computational biology*, 10(11):e1003945, 2014.
- [18] W. Poundstone. *Prisoner’s Dilemma*. Doubleday, 1992.
- [19] D. G. Rand and M. A. Nowak. Human cooperation. *Trends in cognitive sciences*, 17(8):413–425, 2013.
- [20] D. G. Rand, M. A. Nowak, J. H. Fowler, and N. A. Christakis. Static network structure can stabilize human cooperation. *Proceedings of the National Academy of Sciences*, 111(48):17093–17098, 2014.
- [21] R. L. Riolo, M. D. Cohen, and R. Axelrod. Evolution of cooperation without reciprocity. *Nature*, 414:441, 11 2001.
- [22] H. Rusch. The evolutionary interplay of intergroup conflict and altruism in humans: a review of parochial altruism theory and prospects for its extension. *Proc. R. Soc. B*, 281, 2014.
- [23] F. C. Santos, J. M. Pacheco, and T. Lenaerts. Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proc Natl Acad Sci U S A*, 103(9):3490–3494, 2006.
- [24] F. C. Santos, J. M. Pacheco, and B. Skyrms. Co-evolution of pre-play signaling and cooperation. *J Theor Biol*, 274(1):30–35, 2011.
- [25] F. C. Santos, M. D. Santos, and J. M. Pacheco. Social diversity promotes the emergence of cooperation in public goods games. *nature*, 454(7201):213–216, 2008.
- [26] F. P. Santos, F. C. Santos, F. S. Melo, A. Paiva, and J. M. Pacheco. Dynamics of fairness in groups of autonomous learning agents. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2016.
- [27] F. P. Santos, F. C. Santos, and J. M. Pacheco. Social norms of cooperation in small-scale societies. *PLoS computational biology*, 12(1):e1004709, 2016.
- [28] K. Sigmund, C. Hauert, and M. A. Nowak. Reward and punishment. *Proc Natl Acad Sci U S A*, 98(19):10757–10762, 2001.
- [29] B. Skyrms. *The stag hunt and the evolution of social structure*. 2003.
- [30] B. Skyrms. *Signals: Evolution, Learning & Information*. Oxford Univ Press, 2010.
- [31] G. Szabó and G. Fáth. Evolutionary games on graphs. *Phys. Rep.*, 446(4-6):97–216, 2007.
- [32] A. Traulsen. Mechanisms for similarity based cooperation. *European Physical Journal B*, 2008.
- [33] A. Traulsen and M. A. Nowak. Evolution of cooperation by multilevel selection. *Proc Natl Acad Sci U S A*, 103(29):10952–10955, 2006.
- [34] A. Traulsen and M. A. Nowak. Chromodynamics of cooperation in finite populations. *PLoS ONE*, 2007.
- [35] R. L. Trivers. The evolution of reciprocal altruism. *The Quarterly review of biology*, 46(1):35–57, 1971.
- [36] S. Van Segbroeck, J. M. Pacheco, T. Lenaerts, and F. C. Santos. Emergence of fairness in repeated group interactions. *Physical review letters*, 108(15):158104, 2012.
- [37] D. S. Wilson. A theory of group selection. *Proc Natl Acad Sci U S A*, 72(1):143–146, 1975.