Effective Games in Multilayer Complex Networks

Miguel Dziergwa de Carvalho
miguel.dziergwa@ist.utl.pt

Instituto Superior Técnico, Lisboa, Portugal

September 2018

Abstract

Understanding the mechanisms behind the origin and maintenance of cooperation has been the focus of much research during the last decades. In that context, the underlying structure of social interactions has been shown to greatly influence the chances of reaching high levels of cooperation. However, real networked systems are often shaped by multiple interdependencies that are not conveniently captured by a single network. For instance, in social networks, individuals can be connected through different types of relationships — originating from collaboration, professional, friendship, or family ties — which can be conveniently described by a multilayer network. In this thesis, we analyze how a population structured as a multilayer network of interactions alters the chances of reaching cooperation. In particular, we aim to understand under which conditions networked interactions effectively transform, globally, the social dilemmas of cooperation that individuals locally face. To that end, we implement a novel numerical tool that allows us to track the self-organization of cooperation in networked populations with an arbitrary number of layers. Our results show that interactions in multiple layers can transform the original dilemmas, creating new basins of attraction and stable equilibria, absent in a single type of network. Finally, we show that these game transformations are not trivial. Cooperation may either increase or decrease, depending on factors such as the number of layers, the strength of the dilemma, the topology of the network, or the level of degree overlap among layers.

Keywords: Cooperation, Multilayer, Network, Transformations, Game Theory

1. Introduction

Although Network Science [7, 1] is a relatively new field of study – which combines tools and principles from physics, applied mathematics and computer science – it has proven helpful to several other areas of research. The study of interactions and the topology of networks presents itself as a map to interpret phenomena in Economic [8, 9], Ecological [11], and Sociological [10] Systems, as well as in Cognitive Sciences [4].

In social systems, and from a purely theoretical perspective, individuals can be abstracted as nodes and the relations between them as links in a network. The links can represent, for instance, an acquaintance relation, a co-working connection or simply that individuals share a common space such as a gym. After having a suitable representation, one can describe any particular rules that dictate how individuals interact and learn new traits or strategies through those links.

A popular approach to model interactions between agents, and the payoffs herein accrued, is to use Game Theory [6]. In this context, Game Theory allows the modeling of real-life situations while providing an intuitive setting. In particular, we shall use Game Theory to model cooperation in social dilemmas. These are scenarios in which interactions may represent a contest over a shared resource or a conflict between individual and social interests. Different cooperation dilemmas can be obtained by conveniently tuning a few game parameters.

Empirical observations often suggest that the observed behavior of the population does not match the expected outcome derived from a Game Theoretical rational. Cooperation represents an obvious example; it is widespread in nature and societies, from simple organisms to complex Human interactions, even if purely rational analysis would suggest the opposite [22].

We hypothesize that the differences in the theoretically predicted outcomes in cooperation games, and the observed behaviors in reality, differ as a byproduct of different social networks of interactions. Decision-making in a network may thus alter in a fundamental way the global dynamics of cooperation in a population, an hypothesis supported by both experimental [17, 16] and theoretical [12, 18, 20] frameworks. Ultimately, networks lead to a *game transformation* [24, 14, 15], a break in the symmetry between the local and global dynamics of the population, which one can capture by studying the temporal behavioral evolution of the networked population.

Most previous studies, however, neglected the fact that networked systems are often shaped by multiple interdependencies that are not conveniently captured by a single interaction or social network. Individuals can be connected through different classes of ties, originating from collaboration, professional, friendship, or family ties, leading to a multilayer network [5], where each layer stands for a particular type of relationship between nodes.

In this thesis, we will analyze how simultaneous pairwise strategic interactions in multiple layers – here modeled through multiplex networks – can transform the original two-person dilemmas (e.g. Prisoner's Dilemma), fostering (or not) cooperation when compared with a single layer or network. To do so, we will extend the **AGoS**, a tool introduced by Pinheiro et. al. [14, 13], to a multilayer architecture where we vary the structure of the networks, the degree that nodes share across layers, and the importance given to the relative fitness (intensity of selection) of others.

Finally, the abstract nature of the new computational tool we develop renders this framework readily applicable to other time-dependent processes that may also occur in multilayer networks, such as N-person collective dilemmas, opinion dynamics, ecological processes, or disease spreading, among others.

2. Background

Evolutionary game theory applies principles from evolutionary theory to populations of individuals whose interactions follow a game theoretical setting. Individuals (assumed to be rational decision-makers that look to maximize their rewards) change or adapt their strategies, based on the returns obtained from interacting with others. It is, thus, a mathematical framework that allows us to formalize these interactions, through the definition of the range of possible strategies and the payoffs each strategy has (as a function of the strategies others are using). By letting individuals interact over a sufficient amount of rounds, the best strategies (the ones that return a higher payoff) start to be used by a greater amount of individuals, similar to the way a favorable trait would gain prevalence in the context of a Darwinian natural selection process, while a less fit strategy (with a lower payoff) would disappear.

We will consider pairwise interactions between individuals, so called 2-person games. Each type of game (Prisoner's dilemma, Stag-Hunt game, Snowdrift game) has a payoff matrix associated (the names come from famous examples representing the dynamics the matrix originates) and individuals can either behave as Cooperators (**C**) or Defectors (**D**).

When both individuals cooperate they get a reward *R*, when both defect they get a punishment *P* and when one cooperates and the other defects, the former gets *S* (also known as the sucker's payoff) and the latter gets *T* (also know as temptation to defect). The dilemmas occur when individual choices result in defecting instead of, the more beneficial, mutual cooperation. By normalizing *R* to 1 and *P* to 0 we only have two variables remaining, *T* and *S* (instead of four originally) which will be in the range of $0 \leq T \leq 2$ and $-1 \leq S \leq 1$. This simplification has no loss of generality [18]. We will simplify further by defining one single variable **B** (benefit) (that can be unfolded as T = **B** and S = 1-**B**). This takes advantage of the previous knowledge regarding variations in behavior that occur in Prisoner's dilemmas (our main dilemma of interest) as a function of **B**: the incentives for Defection are greater (for a PD) as T increases and S decreases, which is exactly the change we produce by increasing **B**.

The *replicator equation* [22] describes the evolutionary dynamics of infinite and well-mixed populations. Thus, we describe how the frequency of a strategy i would evolve in a population:

$$\dot{x_i} = x_i[f_i(x) - \phi(x)], \quad \phi(x) = \sum_{j=1}^{n} x_j f_j(x) \tag{1}$$

where x_i is the fraction of type i in the population, $f_i(x)$ is the fitness of type i and $\phi(x)$ is the average population fitness. In the particular case when we only consider 2 strategies the equation above simplifies to:

$$\dot{x} = x(1-x)(f_C - f_D) \tag{2}$$

where x represents the fraction of cooperators ($(1-x)$ the fraction of defectors) and f_C represents the fitness of cooperators (f_D the fitness of defectors). This derivative $\dot{x}$ can also be seen as the gradient of selection $g(x)$. If $g(x) > 0$ then the number of cooperators increases, if $g(x) < 0$ then it decreases. In general, strategies whose fitness exceeds the average fitness $\phi(x)$ tend to spread while others tend to die out. The replicator equation describes evolution for well-mixed populations, and only constitutes a good approximation when modelling infinitely large populations. In reality, populations are finite and potentially small; also, individuals tend to interact following an underlying network of contacts and deviating from the well-mixed assumption, where every pair of individuals is equally likely to interact. This way, we will be working with numerical simulations that assume finite populations, which means we should no longer use the *replicator equation* and have to start considering that agents change strategies stochastically as depicted in figure 1, i.e, agents choose to imitate behavior probabilistically, based on the success of the other strategy.

In particular, we will model a process of social learning (known as the *Fermi update rule* [23]) where at each time-step a randomly chosen individual X imitates one randomly chosen neighbour Y with a probability p that increases with their fitness difference

$$p = [1 + e^{-\beta(f_Y - f_X)}]^{-1}. \quad (3)$$

The value of β is the intensity of selection and f_X (f_Y) is the fitness of agent/node X (Y), the latter being a measure of how well an agent is performing. With these elements we can start a simulation to study the emergent behavior on a macro-, population-wide level, stemming from simple rules governing how individuals interact, at micro-level, [21].

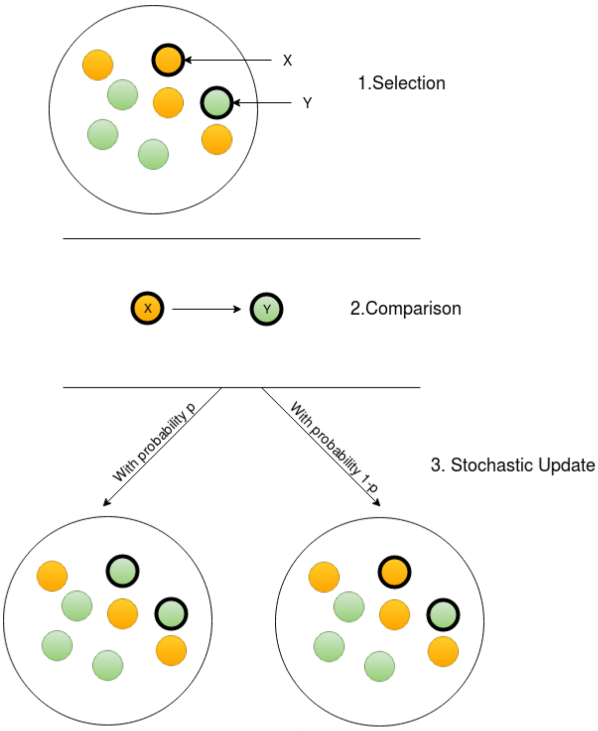

Figure 1: Evolution process in finite well-mixed populations. In networked populations, Y is randomly selected from X's first neighbors.

Furthermore, in well-mixed populations, individuals with the same strategy are considered equivalent; this is not true in structured populations where the context surrounding an individual has to be taken in consideration as well. Individuals with the same strategy are not necessarily equivalent since they also depend on their position in the network and on their neighborhood, at each time step. This motivates the use of the **AGoS**.

The average gradient of selection — **AGoS** — is a time-dependent variable that is used as a tool to study games on graphs. It can provide the same information that the replicator equation does, but in structured populations. The **AGoS** is defined as the average over all possible transitions taking place in every node of the network throughout evolution and the average over many generations. It considers the influence each node can have on its neighbors. The **AGoS** is useful since it captures the population-wide dynamics due to its mean-field character. It can be computed for arbitrary intensity of selection β, for arbitrary finite populations and for any kind of network structure. We will use it to track the self-organization of cooperators when interacting with defectors [14].

We will consider two types of network topologies:

- **Homogeneous random (HR) graphs** [19]: where every individual has the same number of neighbors, but is not necessarily attached to all the nodes in the graph (as in a complete graph), and the neighbors are random (figure 2). In these networks, individuals with the same strategy are not necessarily equally fit because they now have a limited amount of connections, thus they unable to contact with everyone, contrary to what happens in well-mixed populations. This makes the fitness of an individual context-dependent. An example would be a cooperator surrounded by cooperators, or surrounded by defectors, obviously he would survive better in the former scenario. This network is created by using a regular graph and then randomizing the links.

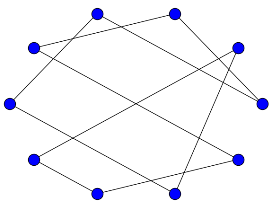

Figure 2: A HR graph with average degree *k=2*.

- **Scale-free (SF) graphs** [3]: where the degree distributions decay with a power-law, showed in Fig. 3. The best way to create a graph of this type is to use the method proposed by Barabási and Albert [3, 1] which combines growth and preferential attachment (we will refer to these as **SFBA**). This method consists in continuously adding nodes to a network, that starts with m_0 nodes, in a way that the probability of linking to an existing node is proportional to a node's degree. This way, nodes with higher degree or hubs, are more likely to attach to the new nodes therefore allowing for the preferential attachment property to emerge in the network.

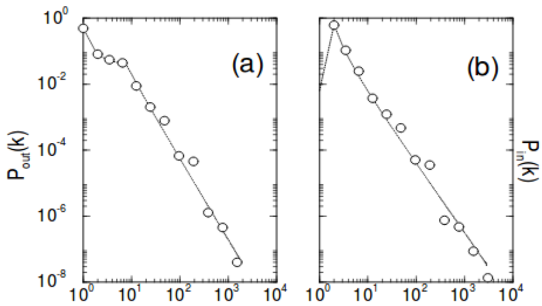

Figure 3: The distribution of (a) outgoing links (URLs found on an HTML document) and (b) incoming links (URLs pointing to a certain HTML document), that follow a power law, meaning we are in the presence of a scale free network,from [2].

We will only use a particular type of multilayer, the *multiplex*, where all the layers have the same set of nodes and the only difference among layers is the way that the nodes are connected among themselves.

3. Methodology

We can divide our process in two separate parts: The first part consists in creating and validating our framework, by replicating the results obtained in previous works (detailed in the dissertation); The second part consists in extending the previous methods to model the evolutionary dynamics of multilayer networks, in particular the interactions and transformations occurring in them. In this section we provide a description of the latter.

We are interested in :

- **Studying the final outcome of cooperation** and how it varies as a function of the game individuals play, reflected by **B**, as well as a function of the number of layers. These results will be interesting because they will raise the questions that we want the **AGoS** to answer. If we observe that the number of layers increases the FFC for, e.g (**B** $= 1.45$), then this is expected to occur due to the interplay between two things: either less cooperators are required to reach a stable fraction of cooperators and, thus, the root of the gradient should be smaller; or the gradient is positive and stronger (in absolute value) and exerts a greater pressure towards cooperation.

- **Extending the AGoS to multilayer networks in order to study and characterize the changes that occur** in the evolutionary dynamics present in 2-person interactions between L layers, where the 2 nodes are randomly selected nodes and the update follows a Fermi function where the fitness of an individual is calculated by averaging over the accumulated payoffs he obtained considering all L layers, instead of the previous payoffs (equation 3) obtained solely in a single layer:

$$p = [1 + e^{-\beta(\frac{1}{L}(\sum_{l=1}^{L} P_y - \sum_{l=1}^{L} P_x))}]^{-1} \qquad (4)$$

where β is the intensity of selection and L is the total number of layers.

In the upcoming sections we will relax the language and use the term *multilayer* often to refer to *multiplex* structures (the different types are described in the dissertation).

We will study how the gradient of selection changes by increasing the amount of layers of the same type, e.g a couple of HR networks. We will perform this analysis for different values of **B**. We will do this for HR networks and SFBA without correlations.

4. Results & discussion

4.1. How multilayer networks impact cooperation

In the following two subsections we finally show our main results. Revisiting, our goal was to understand how does the multilayer aspect impact the overall levels of cooperation, for different networks. On this extended abstract we show the results of considering HR networks and SFBA without correlations, deferring the other case to the dissertation. In figure 4 we show the FFC after 1000 generations obtained for different types of network and for different values of **B**, reflecting the strength of the dilemma.

- In HR networks, we can see two regimes, one on each side of **B** ≈ 1.0125, which we call the **B**$_{inflection}$. We associate the right side, where **B** $>$ **B**$_{inflection}$, with a (what we call) "high" **B**, and the left side, where **B** $<$ **B**$_{inflection}$, with a "low" **B**. When **B** is high, a single-layer can sustain cooperation while multilayers cannot. For low **B**, multilayers can reach higher levels of cooperation.

- In SFBA without correlations, we do not observe an inflection point, meaning that the outcome is correlated linearly with **B**. In this case, for very high or very low **B**, cooperation is constant and independent of the number of layers. For all **B** in between, a single layer provides slightly better conditions for cooperation to emerge than multilayers.

In figure 5 we present a summary of the impact of having $L = 2$ in when comparing to $L = 1$, for the different network topologies.

After observing our results we had the confirmation that the value of **B** was an extremely relevant

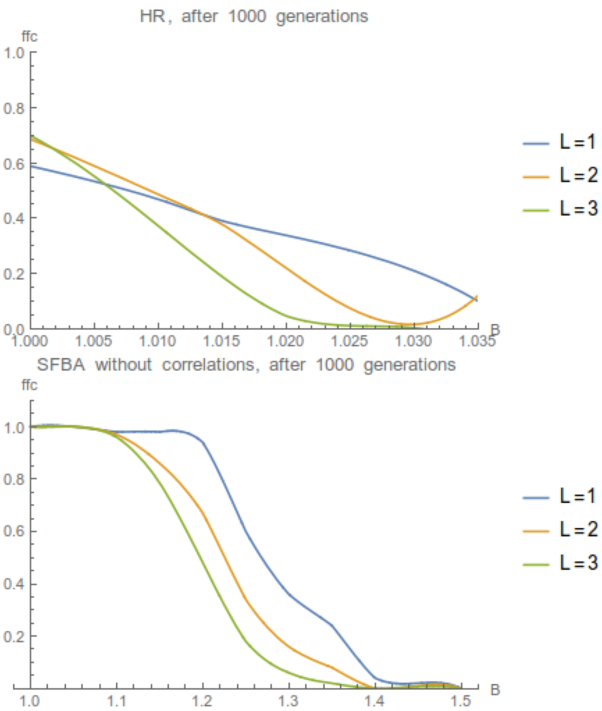

Figure 4: The FFC as a function of **B** in different networks with different layers (**B** $=$ benefit, ffc $=$ average final fraction of cooperators, $L =$ total number of layers)

Network topology	Benefit (B)	Cooperation
HR	High	Decreases
HR	Low	Increases
SFBA(with correlations)	High	Increases
SFBA(with correlations)	Low	Decreases
SFBA(without correlations)	Very high	Maintains
SFBA(without correlations)	Very low	Maintains
SFBA(without correlations)	In between	Decreases

Figure 5: The impact of multilayers in cooperation, for $L = 2$, when in comparison to $L = 1$

variable. In the limit, it would be interesting to apply the multilayer **AGoS** for many different values of **B**. Since we had limited time and computational resources we had to choose some values of **B** that we deemed of greater importance. Therefore we chose, for every network, one value of high and low **B**, where we could observe relevant differences in cooperation, and applied the **AGoS** for each scenario.

4.2. Understanding the changes by analyzing the AGoS

In this section we show the results obtained by calculating the **AGoS** in multilayer networks, during 30 generations (we experimented for a larger number of generations in HR networks and the results were qualitatively the same). By using this tool we aim to understand the impact that multilayers can have in the emergence (or not) of cooperation, by altering the global dynamics of strategy adoption – here quantified through the AGoS – and leading to the results described in the previous section.

- In figure 6 we compare the **AGoS** in HR networks, for two different values of **B**. When considering a low **B**, we can see that multilayers are able to form a basin that goes from $0.5 \leq j/N \leq 0.7$, while the single-layer has a well-defined stable root $x_{stable} \approx 0.5$. The other significant difference is that, for $j/N \geq 0.5$, $|G^A(j, L = 1)| > |G^A(j, L = 2)|$, and since they are both negative, this means that the single-layer exerts a greater pressure towards x_{stable}.

  When considering a high **B**, we can observe that multilayers do not evidence a root, meaning that they cannot withstand cooperation and the game is a defection dominance one. On the other hand, a single-layer shows a stable root $x_{stable} \approx 0.15$, meaning that cooperators can have a chance to survive, due to the existence of this co-existence point.

  In summary, for low **B** the game remains qualitatively the same but with a net increase in cooperation, while for high **B**, it changes from a co-existence to a defection dominance one.

- In SFBA without correlations (figure 7), when we have a low **B** we can see that a single-layer has an unstable root $x_{unstable} \approx 0.3$, which is to the left of all the roots that the multilayer evidences. This means that a single-layer requires less cooperators to reach the tipping point required to ”jump” towards full cooperation. For $j/N > x_{unstable}$ we can observe that a single-layer exerts a larger pressure towards cooperation than a multilayer does. We could argue that multilayers show many roots, but we are not fully prepared to make that claim since these can be attributed to fuzzyness in the simulation.

  For high **B**, we can also observe that the single-layer root is to the left of the multilayer one, as for low **B**. This time, the multilayer exerts a greater pressure towards defection than a single-layer does.

  In summary, for low **B** the game remains a coordination one. For high **B**, it also remains the same but we can imagine that for a **B** slightly greater than the one we used, the game would

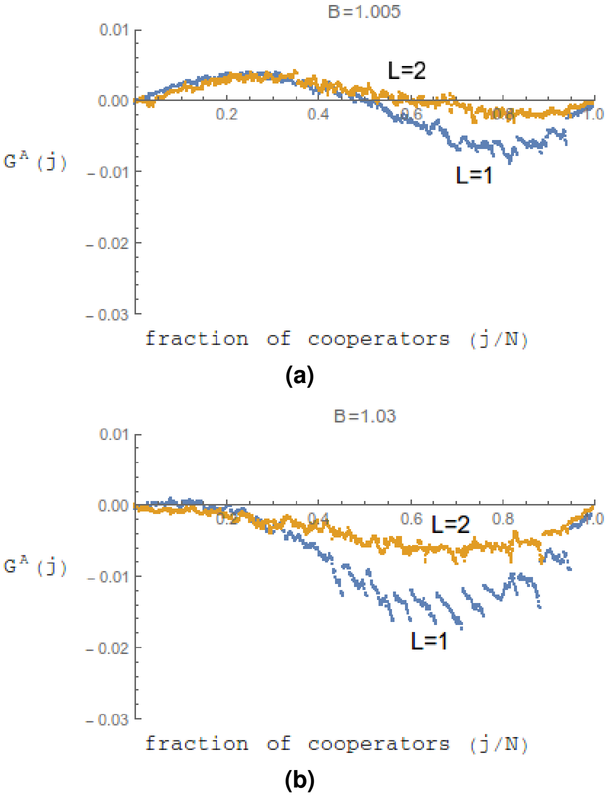

Figure 6: AGoS for HR networks, during 30 generations (**B** $=$ benefit, $L =$ total number of layers, $\beta = 1$)

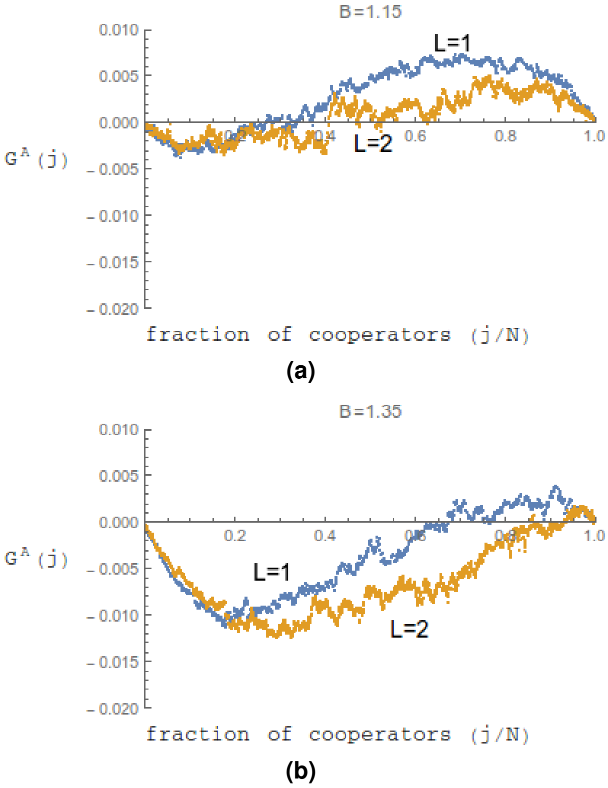

Figure 7: AGoS for SFBA networks without correlations, during 30 generations (**B** $=$ benefit, $L =$ total number of layers, $\beta = 0.1$)

change from a coordination to a full defection one.

After observing the results we are pushed to the conclusion that the multilayer aspect of the networks has a non-trivial impact on the outcome of cooperation. For some topologies, a strong dilemma (characterized by a high **B**) played in a multilayer increases cooperation, while for others topologies it makes it decrease. This change in the number of cooperators sometimes is sufficient for the effective game to change, while other times it is not. This same logic can be observed for weaker dilemmas. In some topologies, the multilayer enhances cooperation and for other inhibits it.

These differences in cooperation can be justified mainly with two arguments. The first is the shift that can occur to the roots of the **AGoS**, and the second is the intensity (or amplitude) of the AGoS.

5. Conclusions

This last section will serve the purpose of summarizing what this thesis achieved and to shed light on some ideas worth pursuing in the continuation of this work.

5.1. Summary of contributions

Throughout this document we always wanted to maintain a rigorous scientific and rational perspective. We did not take anything for certain, and we always questioned what we thought was plainly obvious. This, in part, was the driving force behind wanting to implement a correct framework, validated with as much empirical evidence as we possibly could. This makes for our first contribution: by creating our framework we independently validated numerous previous results, thus solidifying them. The second contribution would be that we achieved results that allow us to observe how multilayer networks affected the emergence of cooperation, in a plethora of network topologies, and for different game intensities (all inside a Prisoner's dilemma). The third and final contribution would be that we managed to explain and justify, by using our multilayer extended version of the **AGoS**, how these changes in behavior occurred inside the networks.

This allows us to claim that we reached our proposed goal to discover what game transformations occur for some different topologies and challenges. We were able to relate the observed FFC with the changes in the **AGoS**. We can also state that the **AGoS** is a tool able to provide useful insight regarding behavior in multilayers, as expected, and can be applied to several other problems.

5.2. Future work

In the end, we look back and can see that we only scratched the surface, if we consider the amount of experiments we can realize and the amount of variables that we could consider. Here we present a profusion of interesting extensions to this work:

- **More network topologies** In this work we only studied HR and SFBA in multiplex networks, yet there are numerous other types with interesting properties that can be used.
- **Mixing topologies** One could also use multilayers that were comprised of different topologies between layers, e.g the interaction between nodes playing in a SFBA and a HR.
- **N-person games** We applied the **AGoS** to 2-person games. It can also be used to study the gradients in games occurring in larger groups.
- **More values of B** We have already argued that the way **B** affects cooperation is far from trivial, therefore a more precise and exhaustive study would perhaps provide useful insight.
- **More layers** We already showed that the multilayers with $L = 2$ have an influence over the outcome. It would be interesting to study the **AGoS** for more layers.
- **Finer degree tuning** We only studied SFBA networks which were either copies or random permutations of one another. It would perhaps be interesting to study them with finer degree correlations.
- **Edge overlap** We applied the changes in degree by switching node *i* with node *j* in the network. It would also be interesting to see how behavior would change if we maintained the degree and randomized a certain amount of the edges of the nodes.

5.3. Final remarks

Generally speaking, the simpler a model, while accurately preserving the key properties of the complex system one wants to address, the "better" it is. This allows for the development of a treatable or transparent model, enough to be comprehended, and later on complexified, if needed. To a great extent, this is what we tried to do here. Cooperation among self-regarding entities is a complex system composed of many interacting components, with an emerging outcome difficult to predict, even in its simplest form, as we have seen. This thesis represents an attempt to resort to such high-level abstractions to characterize the complex interplay between cooperative actions and the network structure that underlies our social interactions. There are, however, aspects that are impossible to address with the present approach. In other words, *less* is not always *more*.

Thus, it is worth pointing out that our model does not capture the metaphysical motifs that are lingering on top of the behavior observed. That is, we can analyze the gradient (or any other hypothetical metric) and see what changes occur and try to justify behavior with them. But we do not think that we are able to get to the driving forces behind decision-making at an individual level, the "what" that moves individuals to act the way they do. We think that there are structures in the human brain that guide our interactions, decisions and the way we experience the world and perceive others, in a way that is far from trivial and that, possibly, is unverifiable, through the study of behavior alone. We believe that this is a realm where biology, psychology, and philosophy, have more to offer than we can.

Despite these doubts, we believe that they are a mere reflection of the limitations of the act of modelling itself, and not the limitations of our work so to speak. These emerge from the fact that we chose a simple update rule that (as any other much more complex rule) is unable to capture to perfection the ever so complex behavior and nature of humans. Yet, if we regard the goals we had, the model proved to be sufficient to allow us to study the behavior on a macro scale, also due to the fact that, even though humans are individuals and can have tremendous variability in many different behavioral aspects, on average we can take a pretty good picture using simple models, which is exactly what we did.

6. Detailed Methods

6.1. Final fraction of cooperators (FFC)

We wanted to see how did the FFC varied as a function of the game parameter, **B**, and the total number of layers, L. We used **B** varying in the interval of $[1, 1.5]$, and L ranging from 1 to 3. We start with 50% of cooperators randomly spread in a population of $L \times 1000$ thousand nodes. Each simulation *p* comprises a fixed pair (**B**,L); we create 100 different multilayer graphs, and run independently for $\Lambda = L \times 1M$ rounds (1000 generations), by the end of which we calculated the FFC for each simulation p using equation:

$$FFC = \frac{1}{L \times P} \sum_{p=1}^{P} \sum_{l=1}^{L} FFC_{l,p} \qquad (5)$$

where $FFC_{l,p}$ is the observed fraction of cooperators, after Λ time-steps, on layer l of simulation p.

6.2. Single-layer AGoS

Our goal is to compute the probability to increase or decrease the number of **C**s ($G(j) = T_A^+(j) - T_A^-(j)$). For every node *i*, the probability of adopting a different strategy at time *t* is:

$$T_i(t) = \frac{1}{k_i} \sum_{m=1}^{\bar{n}_i} [1 + e^{-\beta(P_m(t) - P_i(t))}]^{-1} \quad (6)$$

where k_i is the degree of node *i* and $\bar{n}_i$ is the number of neighbors of *i* that are using a different strategy. Each node has a payoff that resulted from interacting once with each neighbor and we implemented this as playing one game in all edges of the graph. The **AGoS** for a given time *t*, a simulation *p* and for *j* cooperators is:

$$G_p(j,t) = T_A^+(j) - T_A^-(j) \quad (7)$$

where

$$T_A^{\pm}(j,t) = \frac{1}{N} \sum_{i=1}^{AllDs/AllCs} T_i(t) \quad (8)$$

The final and *time-independent* **AGoS** is averaged over all $\Omega = 1000 \times f_i$ simulations and $\Lambda = 30000$ time-steps.

$$G^A(j) = \frac{1}{\Omega\Lambda} \sum_{t=1}^{\Lambda} \sum_{p=1}^{\Omega} G_p(j,t) \quad (9)$$

6.2.1 Multilayer AGoS

Taking in consideration what we described for the **AGoS** in a single layer, in the previous subsection, we will now detail only the differences in calculating them. For $L = 1$, we used $\Lambda = 30000$ time-steps, corresponding to 30 generations. Now, for the number of generations to remain constant for $L > 1$, we have to perform $L \times \Lambda$ time-steps instead. The evolutionary dynamics now considers the differences between the accumulated payoffs, instead of the differences between the payoffs of nodes in a single layer, leading to equation 4.

In order to have the payoffs defined we have to let each node play a game with every neighbor. This is now done by playing a game on every edge of all layers, instead of just one layer. The probability of a node i changing behavior, at time t, is affected by the accumulated the payoffs and is now given by:

$$T_i(t) = \frac{1}{k_i} \sum_{m=1}^{\bar{n}_i} [1 + e^{-\beta(\frac{1}{L}(\sum_{l=1}^{L} P_{m,l}(t) - \sum_{l=1}^{L} P_{i,l}(t)))}]^{-1} \quad (10)$$

where k_i is the degree of node *i* and $\bar{n}_i$ is the number of neighbors of *i* that are using a different strategy. Another change is that now the total pool of nodes has size $L \times N$, leading to:

$$T_A^{\pm}(j,t) = \frac{1}{L \times N} \sum_{i=1}^{AllDs/AllCs} T_i(t) \quad (11)$$

We also implemented a transient of 1 generation where we merely let the population evolve and do not calculate the **AGoS**, to remove any bias. In the end we are still calculating the average gradient for every configuration j/N that the population passes by.

References

[1] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1):47, 2002.

[2] R. Albert, H. Jeong, and A.-L. Barabási. Internet: Diameter of the world-wide web. *nature*, 401(6749):130, 1999.

[3] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

[4] A. Baronchelli, R. Ferrer-i Cancho, R. Pastor-Satorras, N. Chater, and M. H. Christiansen. Networks in cognitive science. *Trends in cognitive sciences*, 17(7):348–360, 2013.

[5] S. Boccaletti, G. Bianconi, R. Criado, C. I. Del Genio, J. Gómez-Gardenes, M. Romance, I. Sendina-Nadal, Z. Wang, and M. Zanin. The structure and dynamics of multilayer networks. *Physics Reports*, 544(1):1–122, 2014.

[6] C. F. Camerer. *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press, 2011.

[7] S. N. Dorogovtsev. *Lectures on complex networks*, volume 24. Oxford University Press Oxford, 2010.

[8] Y. Durham, J. Hirshleifer, and V. L. Smith. Do the rich get richer and the poor poorer? experimental tests of a model of power. *The American Economic Review*, 88(4):970–983, 1998.

[9] C. A. Hidalgo, B. Klinger, A.-L. Barabási, and R. Hausmann. The product space conditions the development of nations. *Science*, 317(5837):482–487, 2007.

[10] M. O. Jackson. *Social and economic networks*. Princeton university press, 2010.

[11] R. M. May, S. A. Levin, and G. Sugihara. Complex systems: Ecology for bankers. *Nature*, 451(7181):893–895, 2008.

[12] M. A. Nowak and R. M. May. Evolutionary games and spatial chaos. *Nature*, 359(6398):826, 1992.

[13] F. L. Pinheiro. Characterization of self-organization processes in complex networks. *PhD Thesis*, 2016.

[14] F. L. Pinheiro, J. M. Pacheco, and F. C. Santos. From local to global dilemmas in social networks. *PloS ONE*, 7(2):e32114, 2012.

[15] F. L. Pinheiro, F. C. Santos, and J. M. Pacheco. Linking individual and collective behavior in adaptive social networks. *Phys Rev Lett*, 116(12):128702, 2016.

[16] D. G. Rand, S. Arbesman, and N. A. Christakis. Dynamic social networks promote cooperation in experiments with humans. *Proc. Natl. Acad. Sci. USA*, 108(48):19193–19198, 2011.

[17] D. G. Rand, M. A. Nowak, J. H. Fowler, and N. A. Christakis. Static network structure can stabilize human cooperation. *Proc. Natl. Acad. Sci. USA*, 111(48):17093–17098, 2014.

[18] F. C. Santos, J. M. Pacheco, and T. Lenaerts. Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proc. Natl. Acad. Sci. USA*, 103(9):3490–3494, 2006.

[19] F. C. Santos, J. F. Rodrigues, and J. M. Pacheco. Epidemic spreading and cooperation dynamics on homogeneous small-world networks. *Phys Rev E*, 72(5):056128, 2005.

[20] 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.

[21] T. C. Schelling. *Micromotives and macrobehavior*. WW Norton & Company, 2006.

[22] K. Sigmund. *The calculus of selfishness*. Princeton University Press, 2010.

[23] A. Traulsen, M. A. Nowak, and J. M. Pacheco. Stochastic dynamics of invasion and fixation. *Physical Review E*, 74(1):011909, 2006.

[24] S. Van Segbroeck, F. C. Santos, T. Lenaerts, and J. M. Pacheco. Selection pressure transforms the nature of social dilemmas in adaptive networks. *New Journal of physics*, 13(1):013007, 2011.