

Sort out Your Neighbourhood: Public Good Games on Dynamic Networks

Kai Spiekermann

Abstract

Axelrod (1984) and others explain how cooperation can emerge in repeated 2-person prisoner's dilemmas. But in public good games with anonymous contributions, we expect a breakdown of cooperation because direct reciprocity fails. However, if agents are situated in a social network determining which agents interact, and if they can influence the network, then cooperation can be a viable strategy. Social networks are modelled as graphs. Agents play public good games with their neighbours. After each game, they can terminate connections to others, and new connections are created. Cooperative agents do well because they manage to cluster with cooperators and avoid defectors. Computer simulations demonstrate that group formation and exclusion are powerful mechanisms to promote cooperation in dilemma situations. This explains why social dilemmas can often be solved if agents can choose with whom they interact.

More than 20 years after Axelrod's seminal computer tournaments and the discussion about direct reciprocity in repeated games [Axelrod1984], we still haven't understood all mechanisms leading to sustained cooperation. Most settings researchers have looked at are based on two unrealistic assumptions. Firstly, most models to explain cooperation comprise games with two players. 2-person prisoner's dilemmas, in particular, have received much attention. But more realistic settings involve more than two persons. Cooperation becomes much harder in multi-person dilemmas. Under realistic assumptions, it breaks down because targeted reciprocation against defectors is not possible. Secondly, most models do not take the effects of social structure into account. They assume random matching of strategies or a tournament of pairwise matching. But the reality of human interaction looks quite different: People are situated in a social network. They meet some people much more often than others. Moreover, people can influence with whom they interact, thereby changing the social structure that determines who interacts with whom.

On the one hand, moving from 2-person to multi-person games makes the emergence and maintenance of cooperation harder. On the other hand, social structure often makes it easier. When we allow agents to change the social network, cooperation can emerge, even in settings quite hostile to cooperation. In this paper I present a model where agents are situated in a network they are able to change over time. They play public good games with their neighbours,

and they cannot observe who of their neighbours defect. This is not an environment where one would expect cooperators to thrive. Nevertheless, simulations show that cooperators can do well because they change the social structure of interaction over time.

This paper is in 6 parts. In section 1 I introduce dynamic networks and discuss how the notion of social structure has influenced recent models of cooperation. This leads to a preliminary, simple model where 2-person prisoner's dilemmas are played on a dynamic network in section 2. Section 3 presents the core model and the simulation results. I discuss the robustness of the model in section 4. Section 5 extends the robustness analysis by taking scale-free network topologies into account. Section 6 draws conclusions.

1 Dynamic Networks

Axelrod himself noted that the feasibility of cooperation increases when the encounter of strategies is positively correlated, that is when cooperators are more likely to meet cooperators [Axelrod1984, pp. 63–69; 158–168]. He considered “clustering” and “territoriality” as possible solutions to the problem that cooperative strategies cannot invade a population of defectors individually. At about the same time, and referring to earlier work from [Axelrod & Hamilton1982], [Eshel & Cavalli-Sforza1982] discussed “assortment” as an important factor to explain how cooperation can be initiated in an evolutionary process when the default behaviour is defection. [Hirshleifer & Rasmusen1989] introduced the idea of ostracism into the economics literature, but they did not explicitly discuss the spatial dimension. In a first wave of literature, with contributions from mathematicians, computer scientists, theoretical biologists, economists, and from other fields, researchers demonstrated the importance of spatial structure for the emergence of cooperation. [Lindgren and Nordahl1994a] discuss iterated prisoner's dilemmas on a lattice to model cooperation in biological systems (see also [Lindgren & Nordahl1994]). Cellular automata with different strategies evolve and spatial structure supports the coexistence of cooperative and non-cooperative strategies. A similar approach by [Ashlock et al.1996] leads to the conclusion that partner selection is an important mechanism for the emergence and stability of cooperation in spatially structured ecologies. [Nowak et al.1994] also analyse the success of cooperative strategies if evolutionary games are played on lattices. Of particular interest for my own approach are papers that incorporate dynamic spatial structures. [Tsfatsion1997] analyses the formation of trade networks where trade interactions resemble iterated prisoner's dilemmas and agents can accept or refuse trading partners. Where partner selection takes place, payoffs are higher, compared to random partner matching.

More recently, a second wave of literature on the connection between spatial structure and cooperation has emerged. This second wave is informed by recent advances in network theory (see for example [Strogatz2001]). Influential papers on the formation of networks in general are [Skyrms & Pemantle2000,

Bala & Goyal2000]. Philosophical applications regarding the problem of cooperation on networks can be found in [Alexander2003, Alexander2007] and [Vanderschraaf2006, Vanderschraaf2007]. Both authors argue that social structure is crucial to explain human cooperation in cooperation dilemmas. Many new contributions in theoretical biology systematically explore the impact of different network topologies [Lieberman et al.2005, Ohtsuki & Nowak2006, Ohtsuki & Nowak2006a, Ohtsuki et al.2006, Santos & Pacheco2006, Hauert & Szabo2003, and further references in these papers]. Of particular interest for this paper is [Santos et al.2008] because they model public good games on networks.

In the aforementioned models the structure of the network remains static. The network structure influences the agents' behaviour and payoffs, but agents are not able to change the structure. This paper, by contrast, implements dynamic network structures, similar to [Pacheco et al.2006] and [Pacheco et al.2008]. The latter offer an analytical treatment of the co-evolution of network structure for a system of direct reciprocity. [Zimmerman et al.2004] have a related discussion for 2-person prisoner's dilemmas. In contrast to these papers, my analysis is concerned with public good games on dynamic networks. Agents can influence the agents they have contact with and thereby shape their neighbourhood. This mirrors the nature of social structure in reality: We have some, but not complete control over the set of people we interact with. We can cut ties with those who cheat us and establish ties with those who seem trustworthy. Such networks can be of a professional (trade networks, academic collaborations, etc.) or a private nature (networks of acquaintance, social networks in virtual worlds, etc.). The approach taken here assumes that the change of network structure happens fast, while the strategies of agents remain unchanged, ruling out the co-evolution of structure and strategy.

Social structure regulates who interacts with whom, and it provides opportunities for agents to change their interaction partners. One typical way to implement the notion of structure is to use grids or—less technically—checkerboards. Each agent inhabits one field on a checkerboard and has a limited number of neighbours. The disadvantage of modeling social structure as a checkerboard is its rigidity: Every field on the board has a fixed number of neighbours in its immediate local neighbourhood. Real social networks look different: Firstly, agents can differ in their number of social contacts; secondly, these contacts are not necessarily local (think of online communities); and thirdly, real agents have the chance to influence the network structure by making and breaking social relations. To incorporate these properties of real social networks, I model social structure as a graph.

A graph consists of vertices and edges. When drawing a graph, vertices are represented as points, and edges as lines connecting these points. Each edge connects two vertices. I take a vertex to represent an agent, and an edge to represent a social relation between two agents. The network in the model is dynamic. It changes its structure because agents can choose to delete edges and new edges are created. This represents the fact that agents have a choice with whom they have social relations.

Analytical solutions to repeated games on dynamic networks are difficult

to find. The space of possible strategies is enormous and the relation between network structure and game strategies is difficult to capture. If there is a large number of agents and a large number of rounds, it is almost impossible to derive an extended game form and “solve” the game. In any case, it is quite implausible to assume that agents have common knowledge of the complete history of game outcomes and network topology. Also, the complex dynamics of repeated games in networks, especially if multi-person games are played, are difficult to tackle analytically. The upshot is that an analytical solution to these complex games is practically impossible and would have to rest on knowledge and rationality assumptions that would render the model unrealistic. Still, these games can be analysed in greater detail. For complex and dynamic games we need to replace deductive analysis with computational modelling.

2 2-Person Prisoner’s Dilemmas

I start with a very simple model to introduce the modelling approach. In the beginning, agents are situated in a social network, with random social structure, constrained by a fixed number of edges (I use different initial network topologies in section 5). For instance, in figure 1, panel a, we see a randomly structured network with 50 agents, with 25 cooperators (white) and 25 defectors (black), connected by 100 randomly drawn edges. Each pair of connected agents plays a 2-person game with payoffs as stated in table 1.

Table 1: Game form with prisoner’s dilemma payoffs.

	cooperate	defect
cooperate	2, 2	-1, 3
defect	3, -1	0, 0

For payoff maximising agents these payoffs constitute a prisoner’s dilemma.¹ However, not all agents are immediate payoff maximisers in this model: I assume that “cooperators” always cooperate, even though cooperation is not a Nash equilibrium for payoff maximisers. “Defectors”, by contrast, always defect. Defectors do better than cooperators in each single game in terms of payoff. However, agents are allowed to sever ties. Deleted ties are replaced by new random ties. Agents can try to sever ties with defectors, hoping that they get connected to cooperators instead. Cooperators aim to cluster with their own kind and avoid defectors, defectors aim to connect to cooperators to exploit them.

More technically speaking, a network is represented by a graph with n vertices and k edges. Let the edges be non-directional. Self-loops (a vertex connected by an edge to itself) are ruled out. Each vertex represents an agent.

¹For simplicity, I will occasionally call a game with the *payoffs of a prisoner’s dilemma* a prisoner’s dilemma, even though cooperators do not play a prisoner’s dilemma in their own perception, all things considered.

Vertices can be of two types: Cooperators (C) and defectors (D). The edges represent interaction relations between agents such that two connected agents interact with each other in each round of the game. In the beginning, the edges randomly connect vertices.² The type of each vertex is also determined at random with the condition that there be x cooperators and y defectors.

In each round, every pair of agents connected by an edge plays a prisoner’s dilemma (in terms of monetary payoff).³ Cooperators always cooperate, defectors always defect. The payoffs for the prisoner’s dilemma are as stated in table 1. After playing the game agents can choose to delete one of “their” edges, i.e. they can choose to delete one of the ties⁴ connecting them to other agents. They can also choose not to delete any edge. This means if an agent i has d edges, i has $d + 1$ alternatives: Delete one of the d edges, or delete no edge.

Different deletion strategies are conceivable. I explore two simple strategies: *zealous* and *inert*. In the first simulation I assume that cooperators are *zealous*. This means they sever ties with defectors whenever they can. Defectors are *inert*, i.e. they never delete connections to other agents, because they benefit from having ties with cooperators and are not harmed by other defectors. After all agents had the option to delete one edge, the number of deleted edges is replaced by new random edges.⁵ This procedure is repeated for many rounds.

Figure 1 shows the effect of repeated play. White vertices represent cooperators, black defectors. In the beginning (panel a) players are randomly connected. After 100 rounds (panel b), the network has changed its structure. Cooperators are only connected to other cooperators, defectors only to defectors. The situation depicted in 1b is stable with the strategies described. Neither cooperators nor defectors have reason to sever any ties, given the strategies *zealous* for cooperators and *inert* for defectors. Since new ties are only established when old ties are deleted, no change in the network structure takes place once cooperators and defectors are completely separated. The payoffs for defectors are higher than for cooperators in the beginning, but separation of the two soon puts cooperators in a better position. In figure 1b defectors receive zero payoff, while cooperators receive payoff 2 for each tie to another cooperator.

With a slight modification the effect becomes even more dramatic. Assume that defectors are *zealous*, too, i.e. they sever ties to other defectors if they can. Figure 2 shows the result. Since defectors no longer keep their edges to other defectors, the cooperators get all edges in the network, and defectors are isolated with no ties to other agents. I have included pseudo-code for this model in the appendix.

Even though the model is simple, it already provides some useful insights.


²Note that the graph is not a complete graph, i.e. typically many pairs of vertices are not directly connected.

³If there are multiple edges between two agents, they play the game as often as there are edges between them. Multiple edges can be interpreted as representing a particularly intensive interaction.

⁴I use the terms “edge” and “tie” interchangeably throughout the paper.

⁵For simplicity, I assume that the new random edge can also be the old, deleted edge. This, of course, is very unlikely for a sufficiently large network. When adding edges, multiple edges between the same agents are allowed.

Figure 1: Complete assortation for 2-person prisoner’s dilemmas with cooperators severing ties to defectors. Cooperators are white, defectors black. (a) is the initial setting with 50 vertices and 100 edges, (b) the network structure after 100 rounds.




Firstly, it shows that network dynamics are a powerful mechanism to enforce cooperation. Without network dynamics, the best cooperators can do is to play a conditional strategy like Axelrod’s TIT-FOR-TAT. Such strategies ‘punish’ defectors with defection. These punishments are costly. By contrast, moving away is a cheap but highly effective punishment, because it imposes future losses on the defector, while giving the punisher a chance to increase payoffs by finding a better partner. Secondly, despite its simplicity, the model gives us a good idea of how some social interactions work. Buyer-seller relations often resemble 2-person prisoner’s dilemmas: The buyer can refuse to pay, the seller can refuse to send the goods (or send faulty goods). If an agent finds that her business partner has cheated her, she stops dealing with him and finds new partners (similar results were reported by [Tesfatsion1997]). In this way business networks of reliable traders emerge even though other enforcement mechanisms are missing (proceeding against someone in a different country is often not worth the effort). However, this will only work if both sides expect future interactions. Without a shadow of the future, neither side has an incentive to cooperate.

3 Multi-Person Public Good Games

More interesting questions arise when cutting ties to defectors is not that easy. Many real social dilemmas involve more than two persons. The paradigmatic cases are collective action and public good problems. When many persons are involved, it is often difficult or impossible to determine who has cooperated or defected. People can get away with free-riding, because there are no effective ways to monitor behaviour and punish defectors. The more anonymous interactions are, the easier free-riding gets. For instance, it is often convenient to

Figure 2: Complete assortation for 2-person prisoner’s dilemmas with cooperators *and* defectors severing ties to defectors. (a) is the initial setting with 50 vertices and 100 edges, (b) the network structure after 100 rounds.



dump one’s rubbish into the street in a moonless night (defect), rather than separating it and carrying it to the next recycling center (cooperate). If no one is watching, or if people do not know each other well enough to identify offenders (think of large anonymous blocks of flats), free-riding remains undetected or unpunished. Therefore, I assume that the behaviour of other agents is not directly observable, i. e. contributions are anonymous. This means that agents only know how many players play and how the outcome differs from the ideal outcome of universal cooperation. They do not learn *who* has defected, unless this can be inferred indirectly. Rather than cutting specific ties to defectors, cooperators can only try to gradually “move away” if they are caught in a neighbourhood with high levels of defection. Surprisingly, cooperators can do well even in public good games with anonymous contributions. People observe how well the production of public goods is going on the aggregate level, and change the social network accordingly. This is the idea for the next model.

I describe the core features of the model here; pseudo-code is provided in the appendix, and some details of the model are explained in the notes. The model is initialised by creating a graph with vertices and edges. Vertices represent agents, edges relations between agents. Each round in the model has three stages: a playing stage, an edge deletion stage, and an edge replacement stage.

I begin by describing the playing stage. There are two types of agents, cooperators (or contributors) and defectors (free-riders). Agents are again represented by vertices in a network. Following Alexander’s (2007) terminology, let the associated group of an agent i be all directly connected neighbours of i and i himself. If a neighbour is multiply connected to i the neighbour features in the associated group as often as he is connected.⁶ In each round, each agent plays

⁶One could say that the multiply connected agent has a higher stake in the game. Also, when agents delete an edge to another agent with whom they are multiply connected, this


a public good game with the agent's associated group. The associated group of i is denoted H_i , and $|H_i|$ is the cardinality of the associated group. Each agent i makes a contribution $c_i \in \{0, 1\}$ to the public good. For notational simplicity, I use superscripts to indicate which agent initiates the game, and subscripts to denote who receives the payoff. The net payoff p_m^i for each participant m in the game initiated by i is

$$p_m^i = \begin{cases} r \frac{\sum_{k \in H_i} c_k}{|H_i|} - c_m & \text{for } |H_i| \geq 2 \\ 0 & \text{for } |H_i| < 2. \end{cases} \quad (1)$$

The enhancement factor r is a parameter with $1 < r < 2$. For convenience, I assume $r = 1.5$. If all agents contribute, each agent receives a net payoff $r - 1$, provided there are at least 2 players.⁷ Defection is the strictly dominant strategy for payoff maximisers. However, cooperation can be a viable strategy if cooperators manage to play the game only or primarily with other cooperators.

Figure 3 gives an example. Agent m has edges with agents a , b , and c , who again have edges with other agents. Remember that all agents play the public good game with their associated group. Here m plays one game with $\{a, b, c\}$, but m also participates in the games initiated by all direct neighbours. Therefore, m participates in *four* games.

Figure 3: A network constellation. Grey circles are defectors, white circles cooperators.



In general terms, the payoff p_m for an agent m is determined by adding the payoffs from all the public good games m is playing, similar to the model of overlapping neighbourhoods proposed by [Santos et al.2008]. This results in

reduces the number of edges between them by 1, rather than deleting all edges.

⁷If agent m is c times (multiply) connected to i , m plays the game as often as he is connected and his net payoff is cp_m^i . This means he contributes c times and gets the payoffs for playing c times.

$$p_m = \sum_{i \in H_m} p_m^i. \quad (2)$$

In the edge deletion stage agents can influence their network by severing ties. The edge deletion proceeds asynchronously in random order. Each agent has the opportunity to delete one edge to one neighbour, or do nothing. After the edge deletion stage is over, the number of deleted edges is replaced by new random edges in the network.⁸

I now turn to possible strategies in the edge deletion stage. Agents are able to gradually change their neighbourhood when the level of cooperation is unsatisfactory. The question is whether cooperators manage to find cooperative neighbourhoods given that they cannot identify defectors directly.⁹ Agents need a criterion to decide if and which ties they should cut. Rational agents should try to determine this criterion by calculating the expected utility gain or loss from severing a tie. To undertake this task it is necessary to understand the information available to an agent after the playing stage is over and before the network change stage begins. The agent does not receive information about who among the neighbours and the neighbours's neighbours is a defector or cooperator, unless this can be inferred from the information described. However, the agent is aware of the network topology in his first and second degree neighbourhood.

After playing in the constellation as shown in figure 3, the following information is available to agent m :

1. Agent m knows the rate of defection in his neighbourhood H_{-m} (H_{-m} is the neighbourhood of m without agent m himself, that is $H_m/\{m\}$).
2. Agent m knows the rates of defection in the groups $H_{a,-m}$, $H_{b,-m}$, and $H_{c,-m}$.

In the example, m knows from the game initiated by m that there is one defector in the immediate neighbourhood. From the other games m knows that a and a 's neighbour are cooperators. He infers that in the set of b and b 's neighbours without m (denoted as $H_{b,-m}$) are 1 defector and 2 cooperators. In addition, m knows that c and his two neighbours are defectors. This in turn also leads to the conclusion that b is a cooperator.

The example demonstrates that agent m is not only interested in the type of the immediate neighbours. Since m is involved in 4 games (one initiated by himself, three others initiated by a , b , and c), m cares about the types of his second degree neighbours as well. In the example, m should delete the edge with agent c , as cooperating with c and her neighbours leads to negative payoffs. However, in general it is not trivial to see whether an agent should sever ties

⁸Again, it is possible that the new random edges replace edges that have just been deleted. Multiple edges are allowed.

⁹Except for the special case that a cooperator has only one neighbour, or is able to make inferences from the games played.

and to which neighbour. To answer this question it is necessary to calculate the expected¹⁰ payoff change from round t to round $t + 1$ caused by severing an edge $\{m, x\}$. When severing a tie to an agent $x \in H_{-m}$ the expected payoff changes in two ways. Firstly, x no longer participates in the game initiated by m . Secondly, m no longer participates in the game initiated by x . The expected payoff change $\Delta E(p_m)$ is

$$\Delta E(p_m) = [p_{m,-x}^m - p_m^m] - p_m^x \quad (3)$$

with $p_{m,-x}^m$ being the payoff m receives from the game initiated by m but played without agent x .¹¹ The value of the term $p_{m,-x}^m - p_m^m$ depends on whether x is a cooperator or a defector. In most cases m does not know the types of his neighbours. However m is able to estimate the probabilities that a neighbour is a defector based on the outcomes of previous games. The complexity of these calculations depends on the sophistication of the agents. With perfect memory and assuming that agents are able to start with suitable prior beliefs, a Bayesian treatment would be possible. A Bayesian agent uses all information that becomes available during the course of the game and updates her beliefs about the types of all other agents. In this paper I restrict myself to much simpler and arguably more realistic heuristics.

There are two reasons why a Bayesian calculation of the expected utility change is of little practical relevance. Firstly, it is unlikely that agents individually or on average behave like perfectly rational Bayesian agents, as the level of computational effort is enormous.¹² Agents have incomplete information and limited cognitive abilities. Therefore they have to use simplifying heuristics to make decisions. Secondly, it is more interesting to show that even rather unsophisticated agents can reach structures where cooperators keep defectors in check. In realistic settings, agents use simple heuristics, and it is of little interpretative interest to model agents as much more sophisticated than they actually are.

To explore the ensuing dynamics, it is a good idea to run computer simulations with some plausible strategies. The estimation heuristic underlying all my strategies assumes that agents base their network choice exclusively on the outcome of the games they have played in the current round. The rule I propose is simple: If an agent wants to sever a tie, the agent severs the tie to the agent whose game had the highest rate of defection in the current round (ties are broken with a random choice among those with the highest rate). In the case of figure 3, this choice is obvious: In the game initiated by m himself, m infers that there is one defector in the immediate neighbourhood. When playing the game

¹⁰The change is expected because it rests on the assumption that the rest of the network remains unchanged, which is usually not the case.

¹¹In the special case of multiple edges, the expected payoff must be calculated such that one of the multiple edges is deleted. This means that x participates in one game less initiated by m , and m participates in one game less initiated by x .

¹²Note that if the game is played over many rounds with a limited number of agents, the evidence gathered in the current game should lead to a revision of all earlier reasoning processes based on earlier evidence. This is computationally very demanding.




Figure 4: Panel a shows complete assortation for a public good game with defectors playing *inert* and cooperators playing *zealous* after 100 rounds. There are 25 cooperators and 25 defectors, connected by 100 edges. Panel b shows the average payoff per agent per round over time for cooperators (dashed) and defectors (solid).

initiated by a , b , and c , m realises that c 's neighbourhood has the highest rate of defection (here, without m it is 100%). m infers that c must be the defector. But even if no certain inference can be drawn, it makes sense to assume that the neighbour with the highest defection rate in the neighbourhood is the most likely defector.¹³

To understand the basic dynamics of the game, I analyse two simple strategies already familiar from the 2-person prisoner's dilemma. The *zealous* strategy means that an agent severs one tie to a neighbour if the agent experienced a non-zero rate of defection in his own neighbourhood. The agent deletes the tie to the neighbour with the highest rate of defection.¹⁴ The *inert* strategy means that agents never sever a tie to another agent. From a myopic perspective (looking only one round ahead), defectors should play *inert*, since they are never harmed by any tie to other agents, and severing ties reduces their chances to exploit cooperators. Cooperators, by contrast, should play a less tolerant strategy, and *zealous* is a strategy trying to get rid of ties to defectors.

Figure 4 shows a typical result when 25 *zealous* cooperators play against 25 *inert* defectors, connected by 100 edges, over 100 rounds. Figure 4a reveals that cooperators and defectors are completely separated after 100 rounds, and given the strategies *zealous* and *inert* this is a stable state, i. e. the network will not change any further. We can see in figure 4b that defectors have higher payoffs in the beginning, but cooperators soon do much better than defectors.

¹³Bear in mind, though, that this reasoning is based on the assumption that m can remember nothing but the last round.

¹⁴If there are several agents with the maximum rate of defection, one of them is picked randomly.

Since defectors do not have any connections to cooperators once the network reaches a stable state, their payoffs in all further rounds will be 0. The assortation procedure has led to mutually beneficial ties between cooperators, while the ties between defectors do not benefit the defectors. I ran this simulation 100 times with different random networks as initial setting, and in each simulation complete separation was reached after 100 rounds. The average payoff for cooperators was 2.05, for defectors 0.16.

Defectors “lose” this game because they always end up without any connection to cooperators, that is with zero payoff for all rounds after the assortation is complete and the network is in stable state. The myopic *inert* strategy led to the complete separation of cooperators from defectors. However, if defectors adopt a less tolerant strategy they might be able to avoid a network stable state with complete separation of cooperators and defectors. Can defectors have a better strategy than *inert* against *zealous*? To describe the development of the network, it is useful to distinguish three different types of edges. CC edges connect cooperator to cooperator, DD defector to defector, and CD cooperator to defector (and vice versa). Thinking about the dynamics of the game analysed so far, it is obvious that defectors should sever DD edges because this creates a chance for new CD edges, i. e. opportunities for the exploitation of cooperators. Therefore, defectors should not accept all connections to fellow defectors. Rather they should keep the dynamics in the network going and try to avoid settlement into a stable state with full assortation. To do this, defectors must delete some or all connections with other defectors. I run simulations where defectors do not accept defectors in their neighbourhood to test this intuition.

Let us assume that that all cooperators and defectors play *zealous*. However, with the given parameters (25 cooperators, 25 defectors, 100 edges, 100 rounds) a stable state with complete separation still emerges in all run simulations. Figure 5 shows a typical result after 100 rounds. The *zealous* strategy has not only led to a complete separation of defectors and cooperators, it also left all defectors without any connection to other agents. The average payoffs for cooperators are much better than for defectors, but defectors tend to hold out for longer with *zealous* and have better payoffs before the separation kicks in, indicating that *zealous* at least delays the settlement into complete separation. After the network is in a stable state, however, *zealous/zealous* leads to higher payoffs for cooperators compared to *zealous/inert*, because there are now more edges between cooperators in the stable state. The graph in 5b suggests that cooperators earn a payoff of 4.5 per round. This is what we should expect: Assume by stipulation that all 25 cooperators have at least one edge when the network is in stable state. This is not an unreasonable assumption, given that the average cooperator has 8 edges with the given parameters in stable state. We can then calculate the average payoff by computing all contributions to the public goods games: Each cooperator contributes in the game initiated by herself. Also, each of the 100 edges induces two further contributions. Therefore, we have 225 contributions. Each contribution yields payoff 0.5 (because everyone cooperates). Thus the average payoff is $225 * 0.5 / 25 = 4.5$.

I ran this simulation 100 times (with different random networks as starting

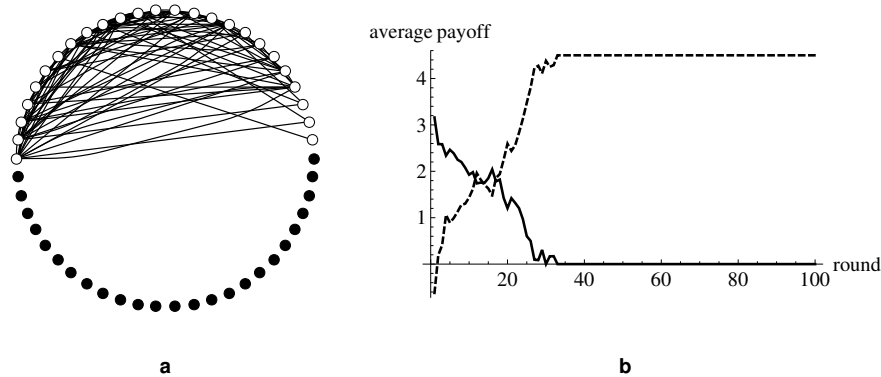


Figure 5: Complete assortment if both cooperators and defectors play *zealous* with 25 cooperators, 25 defectors and 100 edges. Panel b shows the average payoff over time for cooperators (dashed) and defectors (solid).

points), and in all simulations complete separation obtained after 100 rounds. The average payoff for cooperators per round was 3.66, for defectors 0.53.

These results are a powerful demonstration that cooperators can fare very well when they have a chance to cluster. Note that the situation is very hostile to cooperation. Cooperative strategies are usually outcompeted by defectors in repeated anonymous public good games. By providing agents with very limited levels of information, and—crucially—with the option to shape the interaction environment, it is possible for cooperators to cluster and do well.

As pointed out above, the game described is too complex to be analysed analytically, but some observations can still be made. Firstly, with the strategies for cooperators as described above (both *zealous*, or cooperators *zealous* and defectors *inert*) the network will always reach a stable state eventually, even if agents were to use a zero-intelligence strategy for edge deletion and simply chose edges at random. I do not offer a formal proof for this conjecture, but the gist of the argument can easily be seen: There is a small non-zero probability for a transition path from any transient state to a stable state. Therefore the network will eventually end up in a stable state through random drift, certainly in infinite, perhaps in (a very long) finite time. Secondly, the fact that many simulations end up in stable states very fast cannot be explained by random drift alone. The reason for these fast settlements is that cooperators delete CD edges more often and CC edges less often than with the zero-intelligence random edge deletion strategy, thereby pushing the network toward an absorbing stable state. The success of cooperators to reach complete separation from defectors depends on their ability to work towards transitions that are beneficial to them.

The results obtained so far demonstrate that a process of assortment is feasible under specific sets of parameters. Do the specific results hold more generally? Until now the number of agents was assumed to be small, and the

number of edges was limited. Also, it would be important to see how the model behaves if the rate of cooperators to defectors is changed. I turn to these questions in the next section.

4 Parameter Variations and Robustness

To assess the relevance of my model, it is important to show the robustness of its behaviour with different parameter values. A full exploration of the parameter space is infeasible, given the restrictions in computing power and the difficulty to derive analytical results for dynamic models. Nonetheless, it is possible to consider at least some sensible parameter constellations to gain a better understanding of the model's behaviour. I begin with variations in the rate of defectors. I also explore settings with larger networks and networks with increased interconnectivity.

When there are more cooperators than defectors, the network dynamic still leads to complete separation. With 40 cooperators and 10 defectors, both playing *zealous*, and 100 edges in the network, a stable state of complete separation occurred in all 100 simulations after fewer than 100 rounds. Cooperators fared well with an average payoff per round of 2.92, defectors badly (0.14). What happens when there is a rather small group of cooperators playing against a large group of defectors? In 100 simulations with 10 cooperators and 40 defectors, with both types of agents playing *zealous*, no separation occurred after 100 rounds. Figure 6 shows some constellations after 100 rounds. There are still many defectors connected to cooperators, and exploitation of cooperators is widespread. It is also interesting to see that some defectors tend to connect with many cooperators. One can interpret this as a “camouflage” effect. A defector connected to many cooperators is less likely to be identified as a defector. I increased the number of rounds to 1000. The network settled into a stable state in 32 out of 100 simulations. Cooperators experienced low payoffs (0.60), while defectors did better (1.34). For small groups of cooperators it is harder to separate within a reasonable number of rounds.

Larger networks do not differ substantively in their behaviour from the results observed so far. 100 cooperators and 100 defectors, linked by 400 edges, with both types playing *zealous*, behave almost identical to the smaller model: In 20 simulations over 100 rounds, complete separation was always reached and the average payoff for cooperators was 3.80, compared to 0.44 for defectors.

A higher number of edges per vertex can pose a problem for cooperators. In a graph with high connectivity, it is more difficult to distinguish between cooperators and defectors. This is bad for cooperators and good for defectors. Two simulations suggests that it takes more time for the model to settle into a stable state of separation. In simulations with 25 cooperators, 25 defectors, 200 edges, and 100 rounds, with both types playing *zealous*, complete separation was reached in 47 games out of 100. Cooperators had an average payoff of 2.98, defectors 3.67. Longer play turns around the results to the advantage of cooperators: With 1000 rounds, all 100 simulations reached a stable state and




Figure 6: Network constellations after 100 rounds with 10 cooperators and 40 defectors and 100 edges. Both cooperators and defectors play *zealous*.

Table 2: A summary of all simulation results with random initial network and random edge replacement.

coop (x)	def (y)	strategy def	coop, edges (k)	rounds	simu- lations	% stable state	payoff coop	payoff def
25	25	zealous, inert	100	100	100	100	2.05	0.16
25	25	zealous, zealous	100	100	100	100	3.66	0.53
40	10	zealous, zealous	100	100	100	100	2.92	0.14
10	40	zealous, zealous	100	100	100	0	-1.81	1.60
10	40	zealous, zealous	100	1000	100	32	0.60	1.34
100	100	zealous, zealous	400	100	20	100	3.80	0.44
25	25	zealous, zealous	200	100	100	47	2.98	3.67
25	25	zealous, zealous	200	1000	100	100	7.56	0.62

cooperators had an average payoff of 7.56, defectors 0.62. This result of delayed separation is plausible: Since the number of deleted and replaced edges per round does not increase, more rounds are needed to shift the edges. In addition, in a network with large associated groups, the information derived from the defection rates in the neighbourhood has lower quality, compared to networks with smaller associated groups. With the given, limited information it is harder to identify defectors in large groups rather than in small groups. However, in the next section I describe a different edge replacement rule that mitigates this scaling problem regarding connectivity.

Table 2 gives a summary of all my simulations so far. Taking stock, the model displays robustness against most variations. A complete and stable separation of cooperators from defectors is independent of the number of agents. However, if the rate of cooperators is low, separation slows down. Also, a higher connectedness of the network delays the process towards a stable state of separation.

5 Scale-Free Networks and Preferential Attachment

In the last section I tested the robustness of the model under different parameter values. In this section I want to cast the net wider and test the robustness of the result with a different network topology. I distinguish between two different aspects: initial network topology and rules for replacing edges. The simulations considered above were initialised with random networks. Deleted edges were replaced with random edges between any two non-identical agents. These assumptions can be changed, leading to a huge space of possible models. Here I look at least at some plausible assumptions.

Apart from testing the robustness of the results, there are other reasons why different network topologies and rules for changing the topology are worthwhile to consider. While random networks have the merit of theoretical simplicity, they do not seem to be the typical form of real social contact networks. [Newman2003], reviewing empirical results of network research, shows that most social networks have at least two properties: First, many social networks are so-called “small-world networks”; second, the distribution of vertex degrees often follows a power law.

A network is a small-world network if the average path between any two vertices is short, compared to the overall size of the network. The famous (but not necessarily true) hypothesis that any two persons are separated by only “six degrees of separation” is based on the assumption that social contact networks are small-world networks. One simple way to create a small-world network is to start with a “large-world” network, a lattice structure for instance. In a lattice, vertices are only connected to other adjacent vertices. Paths from one point of the lattice to another point can be long. The lattice can be turned into a small-world network by adding a few shortcuts. This reduces the average path length dramatically, because paths between two formerly distant vertices are now shorter due to the shortcuts. Random graphs, as used above, are small-world networks because there are many shortcuts available in the network. Another structural feature that can induce the small-world property is the existence of “superhubs”. A superhub is a vertex with a very high degree (the degree of a vertex is the number of edges connected to a vertex). In a social contact network, a superhub is a person who knows more people than most other persons. For instance, teachers are often superhubs because they interact with hundreds of students. If superhubs exist, the network is usually a small-world network, because each vertex is close to a superhub, and any two vertices can be connected with a short path through the superhub.

Empirical studies suggest that many social networks have a degree distribution that follows a power law in its tail. This is in contrast to random networks, whose degree distribution is binomial. Let p_d be the fraction of vertices with degree d , which is equivalent to the probability of picking a vertex of degree d if one chooses a vertex randomly. If the distribution of probabilities p_d follows a power law, $p_d \sim d^{-\alpha}$, we call the network a scale-free network. The difference

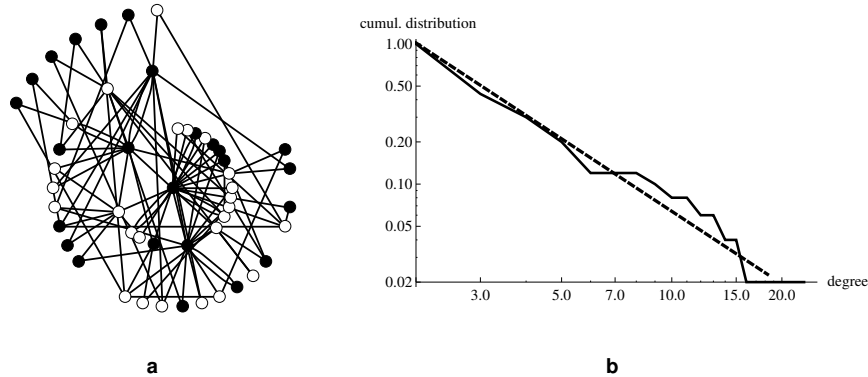


Figure 7: Panel a shows a preferential attachment graph with $m_0 = m = 2$, 50 vertices (25 cooperators, 25 defectors), and 97 edges. Panel b shows the cumulative degree distribution function for this graph on a log-log scale. The line-shaped, downward-sloping curve indicates an approximate power-law degree distribution.

between scale-free networks and random networks is that there are typically more superhubs in a scale-free network than in a random network. Scale-free networks are also small-world networks because of the superhubs.

Why are so many networks scale-free? It is likely that the answer lies in the growth process of networks. If there is a widely shared growth process that produces scale-free networks, this would explain their frequent occurrence. One such growth mechanism is *preferential attachment* [Barabasi & Albert1999]. It proceeds as follows: Start with a small complete graph with m_0 vertices. For each time step, add a vertex and connect it to m existing vertices. For preferential attachment, assume that the probability Π that a new vertex will be connected to an old vertex i depends on the degree of the old vertex, so that $\Pi(d_i) = d_i / \sum_j d_j$. This growth process leads to a scale-free network for large n [Barabasi & Albert1999, p. 511]. Figure 7 shows one preferential attachment graph and its cumulative degree distribution function. One can see a few superhubs and many peripheral vertices. The line-shaped cumulative distribution function indicates that the degree distribution follows approximately a power-law.

Since many social networks are approximately scale-free, it makes sense to test my models on scale-free network topologies for a more realistic setting. I start this exploration by using preferential attachment graphs as initial network topology, but leave the rest of the model unchanged. The graphs are created by starting with 2 vertices, connected by one edge. Then vertices are added, where each connects to 2 existing vertices according to the preferential attachment rule described. The types of vertices are determined randomly such that there are 25 cooperators and 25 defectors. I ran 100 simulations, connected by 97 edges

(1 edge from the starting constellation, 96 edges from the 48 added vertices), all agents playing zealous, with 100 iterations each. Each simulation started with a different preferential attachment graph. All networks quickly reached a complete separation of cooperators and defectors, and cooperators did well (100% stable state, cooperator payoff 3.67, defector payoff 0.47), similar to the results with random initial topologies. This is not surprising: Since deleted edges are replaced by random edges, the network transforms into a near-random network after a few rounds, and therefore the simulation results should be very similar to those with random networks as initial topology.

A more interesting question is how the model behaves when the *replacement* of edges follows a preferential attachment mechanism. As before, deleted edges are replaced by new edges, but this time the random choice of two non-identical vertices is weighted according to the degrees of the existing vertices. A vertex with degree d gets a weight of $d + 1$. I introduce the fixed component 1 to give agents with degree 0 a positive probability to be reconnected again. The probability of a new edge between vertices a and b with degrees d_a and d_b is

$$P(\{a, b\}) = \frac{d_a + 1}{\sum_j (d_j + 1)} \cdot \frac{d_b + 1}{\sum_{j, j \neq a} (d_j + 1)} + \frac{d_b + 1}{\sum_j (d_j + 1)} \cdot \frac{d_a + 1}{\sum_{j, j \neq b} (d_j + 1)}.$$

Call this method preferential edge replacement. It models the phenomenon that agents who already know many people are more likely to meet. This is a more realistic connection mechanism than random connections.

The simulation results confirm the robustness of the model. Again, I start with preferential attachment graphs, but now use the preferential edge replacement mechanism. In 100 simulations with 25 cooperators, 25 defectors, 97 edges, and 100 iterations, the network reached a stable state in all simulations, with high payoffs for cooperators (3.77, compared to 0.37 for defectors). The model is robust regarding the change of edge replacement mechanism. Interestingly, the model with preferential edge replacement also scales better regarding more highly connected networks, compared to non-preferential edge replacement.¹⁵ In 100 simulations with 25 cooperators, 25 defectors, 190 edges ($m_0 = m = 4$), and 100 iterations, a stable state was reached in 79% of the simulations. The average payoffs were 4.78 for cooperators and 2.06 for defectors. This compares to 49% stable states and payoffs 2.71 (cooperators) and 3.57 (defectors) for a random graph model without preferential edge replacement and otherwise similar parameters. For 279 edges ($m_0 = m = 6$), the difference is even more pronounced, as table 3 shows. The preferential edge replacement speeds up the separation of cooperators and defectors, compared to random edge replacement. This shows that cooperators can do well after a few dozen rounds, even in networks with relatively high connectivity. All simulations with preferential attachment graphs are summarised in table 3.

¹⁵I would like to thank one of my anonymous referees for pointing out to me that scale-free networks might scale better in this regard.

Table 3: A summary of simulation results with preferential attachment networks. All simulations are based on the strategy *zealous* for both cooperators and defectors. (PA: preferential attachment).

coop (x)	def (y)	initial network	edge replace	edges (k)	rounds	simul- ations	% stable state	payoff coop	payoff def
25	25	PA	random	97	100	100	100	3.67	0.47
25	25	PA	PA	97	100	100	100	3.77	0.37
25	25	PA	PA	190	100	100	79	4.78	2.06
25	25	random	random	190	100	100	49	2.71	3.57
25	25	PA	PA	279	200	100	90	7.83	2.42
25	25	random	random	279	200	100	35	3.30	5.31
100	100	PA	PA	397	100	20	100	3.83	0.39

6 Conclusion

The agents in this model operate in a setting that is usually rather hostile to cooperation. While repeated 2-person prisoner’s dilemmas can lead to cooperation under suitable conditions, this is not likely in anonymous public good games for two reasons: Firstly, in n -person settings it is not possible to punish specific agents with reciprocal defection. Secondly, if the setting is anonymous, it is not even possible to identify defectors, rendering punishment impossible. The model I propose enables cooperators to do well because it introduces social structure. The key to successful cooperation is a clustering of cooperators and an exclusion of defectors.

The model proposed is simple. It admits only two basic fixed strategies for the playing stage. Future work could consider more advanced strategies that react to previous outcomes. Also, since the playing stage strategies are fixed, the model does not allow for the coevolution of strategies and network structure. This assumption makes sense if strategies are based on dispositions that cannot easily be changed. Nevertheless, it would be interesting to relax this restriction. Moreover, the current model has a fixed number of edges. Several extensions are conceivable where the number of edges changes over time. This would require the introduction of a more sophisticated decision process for the creation and deletion of edges.

The problems of collective action and public good provision have concerned political philosophers and economists for a long time. A lot of energy has been invested into explaining why, empirically, much more cooperation occurs than standard rational choice theory predicts. A focus on repeated games surely points in the right direction, but it only goes half the way. To explain the possibility of cooperation in n -person games with anonymous contributions, one option is to add social structure. This move is particularly attractive because it makes social structure endogenous. Therefore, the model captures an important aspect of social interaction in reality: Social structure and the success of social

interactions are in a dynamic relation with each other. Mutually beneficial interaction reinforces social relations, exploitation weakens them.

This phenomenon is well-known from real settings. For instance, when people venture into joint projects (founding a company, sharing a flat, writing a paper together, etc.), each participant can either contribute or free-ride. It is often difficult to detect free-riding, and even if it can be detected, it is difficult to punish the defector efficiently. Rather, people choose not to continue interaction in groups where the outcome is disappointing. Learning from experience, agents change the social structure of the environment by sticking with groups where free-riding is rare, and staying away from groups where free-riding is common. Individuals willing to cooperate try to cluster in groups with other cooperators, and try to exclude those who defect. What is remarkable about the simulation results is the success of this strategy, even if the information available to the agents is very limited. It is not necessary to track down specific defectors, it suffices to observe the collective outcome and change ties to other agents in response.

I have mentioned mundane examples of cooperation such as flatsharing, co-authoring papers, or running a company as a group of shareholders. However, perhaps the most fundamental problem of cooperation is about life and death: The problem of providing security to live peacefully with each other. In a state of anarchy, where a central provision of policing and security is not possible, security becomes a public good problem. Hobbes reminds us that everyone can kill everyone in a (Hobbesian) state of nature:

“NATURE hath made men so equal in the faculties of body and mind as that, though there be found one man sometimes manifestly stronger in body or of quicker mind than another, yet when all is reckoned together the difference between man and man is not so considerable as that one man can thereupon claim to himself any benefit to which another may not pretend as well as he. For as to the strength of body, the weakest has strength enough to kill the strongest, either by secret machination or by confederacy with others that are in the same danger with himself.” [Hobbes1996[1651], ch. 13]

When government fails, murderers are no longer kept in check by the threat of punishment. Everyone has to fight for himself, and the public good of peace can no longer be provided. In these situations, fleeing to safer areas is often the only option. The large-scale move of refugees in civil wars or failed states can be understood as the attempt to find a safe haven of mutual cooperation in the most basic sense of cooperation: not killing each other. The model discussed in this paper is certainly much too simple to be applied to such complex problems, but with some caveats one could draw the conclusion that overcoming a state of anarchy requires a formation of new local “clusters” of cooperation based on processes of inclusion and exclusion. The model could suggest that states of prolonged anarchy are likely to be followed by a phase of localisation, where

villages or clans form cores of cooperation. Clearly, much more methodological, theoretical, and empirical work is needed to apply computational models to such complex problems. Nevertheless, the example suggests the potential areas of application.

Modelling public good problems with anonymous contributions on dynamic networks shows that cooperation can be maintained, and cooperative agents can do well, if they choose with whom they interact. Cooperators can find each other and build groups of cooperation. Endogenous, dynamic social structure is one important approach to understand the emergence of cooperation.

Pseudo Code for 2-Person Model

The model was coded in Mathematica. I include pseudo code for the all *zealous* strategy for one iteration to describe the central routines of the program used.

```
// INITIALISATION
Create network topology.
Set budget of all agents to 0.
Set status of agents to cooperator or defector.

//MAIN ROUTINE
//GAME
For each edge in network:
  Play prisoner's dilemma with agents connected by edge
  according to their status and change budgets.
End For.

//DELETE EDGES
Set deletedEdges = 0.
For each vertex in network in random order:
  If vertex currently has one or more neighbors
  AND one or more of neighbors is defector then:
    Set cutoffAgent = one of the defectors in neighborhood chosen randomly.
    Change network by deleting one edge from vertex to cutoffAgent.
    Set deletedEdges = deletedEdges + 1.
  End if.
End for.

//REPLACE EDGES
For each edge from 1 to deletedEdges:
  Add random edge (no self-loops) to network.
End for.
```

Pseudo-Code for Public Good Model

This is pseudo-code for one iteration, assuming that both types play strategy *zealous*.

```
// INITIALISATION
Create network topology.
Set budget of all agents to 0.
Set status of agents to cooperator or defector.

//MAIN ROUTINE
//GAME
For each vertex in network:
  Set associatedGroup(vertex) = all agents connected to vertex incl. vertex.
  //agents multiply connected to vertex are counted multiply.
  Play public goods game in associatedGroup and change budgets.
End For.

//RECORD DEFECTION RATES
For each vertex in network:
```

```

neighbors(vertex) = all agents connected to vertex excl. vertex.
For each nb in neighbors(vertex):
  Set otherPlayers = associatedGroup(nb)/{vertex}.
  Set defRates(vertex, nb) = |all defectors in otherPlayers|/|otherPlayers|.
End For.
End For.

//DELETE EDGES
Set deletedEdges = 0.
For each vertex in network in random order:
  If vertex currently has one or more neighbors
  AND |defectors in neighbors(vertex)| > 0 then:
    Set maxdefector = agent with highest rate in defRates(vertex)
    of those currently connected to vertex. //random choice if tied
    If defRates(maxdefector) > 0 then:
      Delete one edge from vertex to maxdefector in network.
      Set deletedEdges = deletedEdges + 1.
    End if.
  End if.
End for.

//REPLACE EDGES
For each edge from 1 to deletedEdges:
  Add random edge (no self-loops) to network.
End for.

```

References

- [Alexander2003] Alexander, J. M. (2003). Random Boolean Networks and Evolutionary Game Theory. *Philosophy of Science*, 70, 1289–1304.
- [Alexander2007] Alexander, J. M. (2007). *The structural evolution of morality*. Cambridge et al.: Cambridge University Press.
- [Ashlock et al.1996] Ashlock, D., Smucker, M. D., Stanley, E. A., and Tesfatsion, L. (1996). Preferential partner selection in an evolutionary study of prisoner’s dilemma. *BioSystems*, 37, 99–125.
- [Axelrod1984] Axelrod, R. M. (1984). *The Evolution of Cooperation*. New York: Basic Books.
- [Axelrod & Hamilton1982] Axelrod, R. M. and Hamilton, W. D. (1982). The Evolution of Cooperation. *Science*, 211, 1390–1396.
- [Bala & Goyal2000] Bala, V. and Goyal, S. (2000). A noncooperative model of network formation. *Econometrica*, 68(5), 1181–1229.
- [Barabasi & Albert1999] Barabasi, A.-L. and Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286, 509–512.

- [Eshel & Cavalli-Sforza1982] Eshel, I. and Cavalli-Sforza, L. L. (1982). Assortment of encounters and evolution of cooperativeness. *PNAS*, 79(4), 1331–1335.
- [Hauert & Szabo2003] Hauert, C. and Szabo, G. (2003). Prisoner’s dilemma and public goods games in different geometries: compulsory versus voluntary interactions. *Complexity*, 4(4), 31–38.
- [Hirshleifer & Rasmusen1989] Hirshleifer, D. and Rasmusen, E. (1989). Cooperation in a repeated prisoners’ dilemma with ostracism. *Journal of Economic Behavior and Organization*, 12, 87–106.
- [Hobbes1996[1651]] Hobbes, T. (1996). *Leviathan*. Rev. ed. Cambridge: Cambridge University Press. First published 1651.
- [Lieberman et al.2005] Lieberman, E., Hauert, C., and Nowak, M. A. (2005). Evolutionary dynamics on graphs. *Nature*, 433(7023), 312–316.
- [Lindgren & Nordahl1994] Lindgren, K. and Nordahl, M. G. (1994). Cooperation and community structure in artificial ecosystems. *Artificial Life*, 1, 15–37.
- [Lindgren and Nordahl1994a] Lindgren, K. and Nordahl, M. G. (1994). Evolutionary dynamics of spatial games. *Physica D*, 75, 292–309.
- [Nowak et al.1994] Nowak, M. A., Bonhoeffer, S., and May R. M. (1994). Spatial games and the maintenance of cooperation. *PNAS*, 91(11), 4877–4881.
- [Newman2003] Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Review*, 45(2), 167–256.
- [Ohtsuki et al.2006] Ohtsuki, H., Hauert, C., Lieberman, E., and Nowak, M. A. (2006). A simple rule for the evolution of cooperation on graphs and social networks. *Nature*, 441, 502–505.
- [Ohtsuki & Nowak2006] Ohtsuki, H. and Nowak, M. A. (2006). Evolutionary games on cycles. *Proceedings of the Royal Society B*, 273, 2249–2256.
- [Ohtsuki & Nowak2006a] Ohtsuki, H. and Nowak, M. A. (2006a). The replicator equation on graphs. *Journal of Theoretical Biology*, 243, 86–97.

- [Pacheco et al.2006] Pacheco, J. M., Traulsen, A., and Nowak, M. A. (2006). Social diversity promotes the emergence of cooperation in public good games. *Nature*, 454, 213–217.
- [Pacheco et al.2008] Pacheco, J. M., Traulsen, A., Ohtsuki, H., and Nowak, M. A. (2008). Repeated games and direct reciprocity under active linking. *Journal of Theoretical Biology*, 250, 723–731.
- [Santos et al.2008] Santos, F. C., Santos, M. D., and Pacheco, J. M. (2008). Social diversity promotes the emergence of cooperation in public goods games. *Nature*, 454, 213–216.
- [Santos & Pacheco2006] Santos, F. C. and Pacheco, J. M. (2006). A new route to the evolution of cooperation. *European Society for Evolutionary Biology*, 19, 726–733.
- [Skyrms & Pemantle2000] Skyrms, V. and Pemantle, R. (2000). A dynamic model of social network formation. *PNAS*, 97(16), 9340–9346.
- [Strogatz2001] Strogatz, S. (2001). Exploring complex networks. *Nature*, 410, 268–276.
- [Tsfatsion1997] Tsfatsion, L. (1997). "A Trade Network Game with Endogenous Partner Selection." In: H. M. Amman, B. Rustem, and A. B. Whinston (eds.), *Computational Approaches to Economic Problems*, Kluwer, 249–269.
- [Vanderschraaf2006] Vanderschraaf, P. (2006). War or peace?: a dynamical analysis of anarchy. *Economics and Philosophy*, 22, 243–279.
- [Vanderschraaf2007] Vanderschraaf, P. (2007). Covenants and reputations. *Synthese*, 157, 167–195.
- [Zimmerman et al.2004] Zimmermann, M. G., J. M., Eguíluz, V. M., and San Miguel, M. (2004). Coevolution of dynamical states and interactions in dynamic networks. *Physical Review E*, 69, 065102(R).