

# Fostering cooperation through dynamic coalition formation and partner switching

Ana Peleteiro, University of Vigo

Juan C. Burguillo, University of Vigo

Josep Ll. Arcos, IIIA-CSIC, Artificial Intelligence Research Institute

Juan A. Rodriguez-Aguilar, IIIA-CSIC, Artificial Intelligence Research Institute

In this paper we tackle the problem of maximizing cooperation among self-interested agents in a resource exchange environment. Our main concern is the design of mechanisms for maximizing cooperation among self-interested agents in a way that their profits increase by exchanging or trading with resources. Although dynamic coalition formation and partner switching (rewiring) have been shown to promote the emergence and maintenance of cooperation for self-interested agents, no prior work in the literature has investigated whether merging both mechanisms exhibits positive synergies that lead to increase cooperation even further. Therefore, we introduce and analyze a novel dynamic coalition formation mechanism, that uses partner switching, to help self-interested agents to increase their profits in a resource exchange environment. Our experiments show the effectiveness of our mechanism at increasing the agents' profits, as well as the emergence of trading as the preferred behavior over different types of complex networks.

Categories and Subject Descriptors: I.2.11 [**Distributed Artificial Intelligence**]: Multi-agent Systems

General Terms: Algorithms, Experimentation

Additional Key Words and Phrases: coalitions, rewiring, cooperation, trading, learning

## ACM Reference Format:

Peleteiro, A., Burguillo, J.C., Arcos, J. Ll., Rodriguez-Aguilar, J.A. 2013, Fostering cooperation through dynamic coalition formation and partner switching. *ACM Trans. Autonom. Adapt. Syst.* V, N, Article A (January YYYY), 32 pages.

DOI = 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

## 1. INTRODUCTION

The global benefit of an agent population is deemed to improve if all agents cooperate. However, achieving cooperation in multi-agents systems (MAS) is a challenging issue, particularly when agents are self-interested [Hogg 1995]. Individual (self-interested) decisions may become detrimental despite eventually delivering short-term benefits. For instance, if many individuals try to download the very same file at the same time, the download speed they will experience is bound to dramatically decrease. Instead, group (social) decisions can result in a mutually beneficial cooperation that holds over time [Shehory and Kraus 1993]. Therefore, designing mechanisms that promote the emergence and maintenance of cooperation for self-interested agents has become a major area of interest in MAS [Doran et al. 1997]. In fact, several coevolutionary game

---

Author's addresses: Peleteiro, A., Burguillo, J.C., Department of Telematics Engineering, Telecommunication Engineering School, University of Vigo; email {apeleteiro, j.c.burguillo}@gti.uvigo.es

Arcos, J. Ll., Rodriguez-Aguilar, J.A., Department of Multi-Agent Systems, Artificial Intelligence Research Institute IIIA-CSIC, Barcelona; email {arcos,jar}@iia.csic.es

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

© YYYY ACM 1556-4665/YYYY/01-ARTA \$10.00

DOI 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

theory approaches have been used to provide a framework to study cooperation in those cases [Perc and Szolnoki 2010].

The emergence of cooperation is often studied in the context of the Prisoner's Dilemma (PD) theoretical framework [Axelrod 1984], which is well-known in game theory. This has been specially useful for understanding the role of local interactions and the maintenance of cooperation [Nowak and May 1992; Langer et al. 2008]. The prisoner's dilemma captures the conflict of interest between what is the best for the individual (defection) and what is best for the group (cooperation), and thus creates a social dilemma [Hogg 1995]. Thus, specific mechanisms for the evolution of cooperation are required to help the population to overcome this dilemma.

To prevent social dilemmas and promote and stabilize cooperation, we distinguish three main strands of work in the literature: coalition-based mechanisms, partner-switching mechanisms, and self-governing institutions. The idea of self governing institutions has its origin on Nobel Laureate Elinor Ostrom [Ostrom 1990]. Works in this line, as Pitt et al. [Pitt et al. 2011; 2012a], use self-governing institutions to model resource allocation in a network applying a formal characterization of Ostrom's socio-economic principles. Coalition-based mechanisms have their roots in the seminal work of Axelrod introduced in [Axelrod 1997] (chapter 6). Axelrod proposes a tribute/tax model that allows agents to achieve cooperation when forming coalitions around some emerging leaders. To maintain coalitions, leaders charge their agents some tribute/tax in favor of some benefit (e.g., guaranteed cooperation, protection against cheaters). This is a clear example of the known tradeoff between the benefits vs. the costs of collaboration (e.g., taxes) [Tanimoto, K. 2002]. Therefore, Axelrod's mechanism is based on a dynamic coalition formation model together with a tax model. Notice that the choice of coalitions comes at no surprise. The notion of a coalition of individuals has been studied by the game theory community for decades, and has proved to be useful in both real-world economic scenarios and multi-agent systems. In fact, coalition formation [Shehory and Kraus 1995; Sandholm et al. 1999] is one of the fundamental approaches in multi-agent systems for establishing collaborations among agents, each with individual objectives and properties. In [Burguillo-Rial 2009] Axelrod's dynamic coalition formation model was successfully adopted to help agents, on grid topologies, cooperate when using a spatial version of the PD [Nowak and May 1992; 1993]. Later on, Salazar et al. [Salazar et al. 2011] introduce an extension of the work in [Burguillo-Rial 2009] that fosters cooperation on complex networks. These networks provide a more realistic model of the topological features found in many nature, social and technological networks (e.g., social networks, the Internet, ecological populations) [Watts and Strogatz 1998; Reka and Barabási 2002]. Furthermore, it is known that they can influence emergence [Pujol et al. 2005]. Salazar et al. design a cooperation mechanism to emerge and sustain full and profitable cooperation, via a single super-coalition, but with a low collaboration cost (tax). Likewise the approaches in [Axelrod 1997; Burguillo-Rial 2009], Salazar et al. also rely on a dynamic coalition formation mechanism, but they also design a consensus mechanism that allows coalition members themselves (instead of leaders) to reach a convention over the fair price to pay to be part of a coalition.

Although the coalition-based mechanisms described in [Axelrod 1997; Burguillo-Rial 2009; Salazar et al. 2011] confirm that coalitions indeed facilitate cooperation between self-interested agents, there is still room for improvement. Firstly, a coalition leader must be paid by the agents belonging to the coalition. Furthermore, notice that a coalition leader imposes her decision on the agents in the coalition to maximize cooperation. Therefore, each coalition leader receives a payoff that is not shared between the members of her coalition. Moreover, the decision-making of each coalition is centralised in a single entity: the leading agent. Additionally, they assume that agents cooper-

ate with their coalition mates. However, assuming cooperation is too restrictive. Other works, as Fu et al. [Fu et al. 2012] and Peleteiro et al. [Peleteiro et al. 2012], consider that an agent should be able to autonomously decide how to behave with respect to coalition-mates and agents outside her coalition. Secondly, although the mechanisms in [Burguillo-Rial 2009; Salazar et al. 2011] promote cooperation on different network topologies, these networks are static (fixed). However, in most real-world situations, the topology of the network changes in response to the state of the network and the other way around, namely the state of the network changes in response to the topology. Research on games on dynamic topologies has found empirical evidence showing that partner switching leads to cooperative behavior. Along this line, Fu et al. [Fu et al. 2009] propose a model of coevolutionary prisoner's dilemma that allows agents to either adjust their strategies or switch their defective partners, showing that partner switching may help stabilize cooperation. Although in a different realm (the investigation of tag-based coordination), Griffiths and Luck [Griffiths and Luck 2010] show that partner switching (rewiring in their terms) can help increase coordination resilience in the face of malicious behavior. However, to the best of our knowledge, no prior work in the literature has investigated whether putting together dynamic coalition formation with partner switching show positive synergies that lead to increase even further cooperation.

Besides the way coalitions are formed, we must also consider why members of a coalition establish cooperation. This is particularly important when considering actual-world scenarios where agents own resources that they can trade. In fact, regarding social and economic collaborating entities (as international alliances, trading agreements, or cooperation among corporations), resource trading plays an important role [Jackson et al. 2003]. Not only in social and economic environments, but we can also consider technological scenarios, as networks of computers that may share their resources, e.g., its CPU, files, etc. To the best of our knowledge, no coalition-based mechanism in the literature has captured the concepts of ownership and trade of resources.

Against this background, our main contribution is a novel mechanism to maximize cooperation among self-interested agents that own resources, where benefits of agents are increased through exchange of their resources. Our cooperation mechanism is based on three main components:

- a game-based interaction model that includes the trading of resources, based on Yee's [Yee 2003] trading model;
- a dynamic coalition formation mechanism that allows agents to: (i) decide whether to join or leave coalitions; and (ii) collectively self-determine the inner and outer behaviors of a coalition (without the intervention of a leader); and
- a partner switching (rewiring)<sup>1</sup> strategy based on experiences acquired in previous interactions that helps agents avoid defective behaviors.

Furthermore, we empirically and thoroughly evaluate our mechanism. We observe that indeed coalition formation plus rewiring allows agents to obtain up to 15% more payoff than only employing either coalition formation or rewiring, and up to 30% more when none of them are employed. However, the benefits of our mechanism depend on the availability of resources, the network topology, and the rewiring frequency employed by agents. Thus, overall our experiments indicate that:

- The higher the availability of resources, the larger the payoff that agents obtain by cooperating, being 40% more when comparing a plentiful resource scenario versus a scarce one.

---

<sup>1</sup>Henceforth we shall employ the term rewiring for shorter.

- Rewiring leads to an increase of gains independently of the topology and its effect in the increase of gains is larger the higher the availability of resources. This increase reaches 20% in a plentiful resource scenario.
- The higher the rewiring frequency, the lower the clustering of the agent population. In other words, the higher the rewiring frequency, the bigger the coalitions formed by agents and the fewer the number of coalitions.
- Agents' strategies adapt to the availability of resources and concrete scenarios to obtain the highest benefits.
- The number of traders increases with the availability of resources, being already more than 80% when the resources are not scarce. Thus trading, i.e., cooperating, emerges as the preferred strategy;

The paper is organized as follows. In Section 2 we thoroughly review the related work in the literature. Next, in Section 3 we introduce our cooperation mechanism, while in Section 4 we offer a detailed empirical analysis. Finally, Section 5 concludes and sets paths to future research.

## 2. RELATED WORK

Individual decisions (self-interested), besides providing only momentary benefits, are detrimental if many agents take them. Instead, group decisions (social) can result in a mutually beneficial cooperation that holds over time [Shehory and Kraus 1993]. This is why coalition formation has attracted the attention of researchers for several years.

In the following sections we present a review of related work in the literature. We begin with a brief review of static coalition formation, which is followed by related work on dynamic coalition formation.

### 2.1. Static Coalition Formation

Coalition formation is a process where agents come together to achieve a goal or to increase their performance. Coalitions have been widely used in multi agent systems [Shehory and Kraus 1993; Shehory et al. 1998; Salazar et al. 2011], as they enable agents to accomplish goals they are unable to accomplish independently. However, the number of potential coalitions increases when increasing the number of agents, since the problem is exponential.

In the following subsections, we present two different approaches to tackle the coalition formation problem. Firstly, optimization approaches, which mainly focus on finding an optimal coalition, i.e. tackling the CSG problem. Secondly, game theoretic approaches, because they have implications and uses in many real-world domains, including those involving automated agents, electronic commerce, auctions, and general resource allocation scenarios.

*2.1.1. Optimisation approaches.* One of the main problems in coalition formation is the coalition structure generation (CSG) problem. It has been shown to be NP-complete, and existing algorithms cannot generate solutions within a reasonable time for even moderate numbers of agents. Thus, to find an optimal coalition can become intractable because the number of coalition structures exponentially depends on the number of agents. There are several algorithms trying to tackle the CSG problem, and according to [Rahwan and Jennings 2008b], they can be classified in three main categories: dynamic programming (DP) [Rahwan and Jennings 2008b; Service and Adams 2011], heuristics [Shehory and Kraus 1998], and anytime optimal algorithms [Sandholm et al. 1999; Rahwan et al. 2009].

Each of those approaches has its advantages over the others, and this led researchers to try to develop new approaches combining their best characteristics. To get the best of DP and anytime algorithms, Rahwan et al. [Rahwan and Jennings 2008a] combine

the state-of-the-art dynamic programming [Rahwan and Jennings 2008b] and state-of-the-art anytime algorithm [Rahwan et al. 2009] on CSG, and develop the IDP-IP hybrid algorithm, which outperforms the time convergence of the former ones. However, the approach presented in [Rahwan and Jennings 2008a] has some limitations, that are identified and solved in [Rahwan et al. 2012]. This work provides a new IDP-IP\* algorithm that outperforms IDP-IP.

Since finding the optimal partition of agents by checking the whole space may be too expensive in terms on time and computation, several other approaches have been proposed to try to reduce the search space and provide faster solutions. Shrot et al. [Shrot et al. 2010] re-examine the computational complexity of the different coalition formation problems when assuming that the number of different agent types is fixed, showing that many of the problems that were proved hard under a general case are polynomial with this assumption. Another option to reduce the search space is to avoid redundant computation, as done in Voice et al. [Voice et al. 2012], where they consider coalition formation problems for agents with an underlying synergistic graph, that constrains the number of feasible coalitions.

*2.1.2. Game theoretic approaches.* One of the goals when forming coalitions is to improve the cooperation among agents. To address this issue, game theory approaches have been widely used. Thus, as a result of the desire to embed game theoretic principles into agent systems, computational aspects of game theory have been extensively studied in recent years [Bachrach and Rosenschein 2008].

One static and theoretical view of coalition formation concerns hedonic games, which are a rich and versatile class of coalition formation games which also encapsulate various stable matching scenarios. The main focus in hedonic games has been on notions of stability for coalition structures such as Nash stability, individual stability, contractual individual stability, or core stability and characterizing conditions under which the set of stable partitions is guaranteed to be non-empty [Aziz et al. 2011]. However, in our work we neither focus on static coalitions nor on finding stability conditions, but in a dynamic environment where agents join, leave and form new coalitions to increase their benefits.

Overlapping coalition formation (OCF) games are cooperative games where players can simultaneously participate in several coalitions. Using this game, Shehory et al. [Shehory and Kraus 1996] present an anytime algorithm that provides suboptimal results. Chalkiadakis et al. [Chalkiadakis et al. 2010] propose models for overlapping coalition formation that allow to handle and reason about stability of overlapping coalition structures. Contrary to these OCF approaches, in our work we consider that agents can only belong to one coalition, as we regard coalitions as independent entities where an agent has to commit exclusively when it joins. This is because in our case coalitions act as a whole and take joint decisions to play against others.

Bachrach et al. [Bachrach and Rosenschein 2008] propose the coalitional skill games (CSGs), a simple model of cooperation among agents to find the optimal coalition structure to improve gains when each agent performs a task. This is a restricted form of coalitional game, where each agent has a set of skills that are required to complete various tasks. In other type of games, the coalitional resource games (CRG), each agent has a set of resources. In order to achieve a set of goals a coalition has to count on the agents that possess the necessary resources. Wooldridge et al. [Wooldridge and Dunne 2006] investigate and classify the computational complexity of a number of natural decision problems for CRGs. However, in our work we focus on agents that have resources that they can trade and all the resources are of the same type. Therefore, our focus is not to find the optimal coalition, but to improve resource sharing in an a priori selfish scenario.

## 2.2. Dynamic Coalition Formation (DCF)

Most previous works in literature aim at finding the optimal coalition. However, when we need to form coalitions in a dynamic environment, where agents constantly may want to change the coalition they belong to, computing the optimal coalition may be either infeasible (because optimality is constrained to a very small number of agents) or take longer than the time of existence of a coalition. Moreover, when the focus is not to find an optimal coalition, or to study the stability of a coalition, the contributions introduced in Sect. 2.1 may not be useful.

Optimal coalition formation has several drawbacks when applied to the real world. First, the time needed to find the optimal coalition prevents its use in a dynamic multi-agent system where agents have to decide if it is beneficial for them to join in a small amount of time. Second, the number of agents involved has to be small, since the number of coalition structures ( $O(n^n)$ ) is so large that it cannot be enumerated for more than a few agents [Sandholm et al. 1999]. Thus, if we focus on a system where we have a large number of agents interacting, the computational cost is so high that it makes it impossible to compute. This is why it is necessary to use domain knowledge and/or mathematical games with certain constraints and where agents have particular characteristics to solve the problem of coalition formation in a reasonable efficient way.

*2.2.1. DCF over static topologies.* To form dynamic coalitions to improve cooperation, there is a need to find decentralized procedures that allow self-interested agents to negotiate the formation of coalitions and division of coalition payoffs. In real world scenarios, it may happen that agents are selfish and only focused on improving their own performance, while if they were cooperative, the whole system performance would improve. To model this situation, non-cooperative games, where agents are selfish, have been used to study coalition formation and its dynamics. Along this line, the Iterated Prisoner's Dilemma (IPD) game has been widely used in modelling various social and economic phenomena, as well as the emergence of cooperation. Contrary to the classic Prisoner's Dilemma (PD), where defecting is the dominant strategy, in repeated games where the total number of rounds is random or unknown, sustained cooperation strategies may emerge [Aumann 1959].

The IPD with coalitions has been firstly used in spatial scenarios, where agents play in a grid scenario, interacting only with the closest neighbors. In this scenario, Seo et al. [Seo et al. 1999] study the emergence of cooperative coalitions in N-player Iterated Prisoner's Dilemma (NIPD), focusing on how the local interaction between agents affects the evolution of the game. They conclude that the more localized the interaction is, the easier is to evolve cooperation. Moreover, to improve their previous results, in [Seo et al. 2000] the authors use the IPD to study coalition emergence in a co-evolutionary learning environment. In this case, they assign to each agent a confidence that specifies how well each is dealing with her opponents. That confidence is adapted through evolutionary learning, basing coalition behavior on most confident agents. With this, Seo et al. found that adaptative confidences can improve coalition performance, and deal with different opponents. Nguyen et al. [Nguyen and Ishida 2009] use the spatial version of the IPD to model one environmental coalition formation problem. In their work, the authors study how specifying different weights for agents, which determine how powerful they are, and allowing them to decide if taking their own decisions or following more powerful agents, influences coalition formation and cooperation.

Also in a spatial scenario, Burguillo [Burguillo-Rial 2009] presents a framework for a memetic analysis of coalition formation considering the spatial prisoner's dilemma. In his approach, agents may play isolated or join coalitions ruled by leaders. Each leader defines the behavior of the agents belonging to her coalition and charges them with

taxes. Using memetic strategies, Burguillo obtained emerging cooperative coalitions. However, Salazar et al. [Salazar et al. 2011] noticed that in Burguillo's work: (i) the coalition strategies employed cannot accomplish full cooperation in complex networks; (ii) the mechanism employed by leaders to tax agents is unfair for the population as a whole. To address these issues, the authors propose a set of strategies that promote profitable cooperation in complex networks, as well as a consensus mechanism to reach a convention to establish fair taxes. Although the coalition-based mechanisms described in [Burguillo-Rial 2009; Salazar et al. 2011] confirm that coalitions indeed facilitate cooperation between self-interested agents, there is still room for improvement. Firstly, in those two approaches, a coalition leader must be paid by the agents belonging to the coalition. This penalises the utility that an agent can obtain from participating in a coalition. Furthermore, a coalition leader imposes her decision on the agents in the coalition to maximize cooperation. By imposing the coalition's strategy, the leader does not take into account valuable information that agents could use for the benefit of all the members of the coalition, not only the leader. Moreover, the payoff distribution is unfair, since: (i) leaders receive a payoff that they do not distribute; and (ii) the agents in the frontier of the coalition (interacting with agents outside of the coalition) obtains less payoff. In order to avoid the mentioned shortcomings, in this paper, we focus on *flat* coalitions, namely coalitions without leaders that: (i) collectively self-determine their behaviors; and (ii) share profits without applying taxes.

When agents group, they must decide how to behave with in-group and out-group agents. Fu et al. [Fu et al. 2012] propose a mathematical framework for the evolution of in-group favoritism, where different agents may use different strategies with agents in its group or outside of it, updating its behavior depending on the payoff. They determine under which conditions in-group cooperation emerges, and when out-group cooperation may be beneficial, using a mutation-selection process. In our case, we also allow agents to decide their behavior against in and outsiders of their coalitions. However, we want agents choose their behavior depending on their own past experience.

Some works have used learning algorithms to address coalition formation and behavior decisions, since these algorithms allow agents to modify their behavior depending on their past experiences. In Bazzan et al. work [Bazzan et al. 2011], agents play the spatial IPD considering social relationships such as belonging to a hierarchy or to a coalition, or being isolated. Each of the members individually learns, using Q-learning, which is the best behavior, whether to be isolated or in a group, and if behaving cooperatively or defectively. Moreover, if they belong to a group, they vote to decide the whole coalition behavior. A preliminary version of this approach was also tested with positive results in [Peleteiro et al. 2011], where the focus was to compare Q-learning and learning automata techniques using different payoff matrix. In both works, results show that learners end up joining in coalitions and cooperating among themselves.

However, previous approaches [Bazzan et al. 2011; Peleteiro et al. 2011] assume that agents behave cooperatively within a coalition. Assuming cooperation is too restrictive, Peleteiro et al. [Peleteiro et al. 2012] consider that an agent should be able to autonomously decide how to behave with respect to coalition-mates and agents outside her coalition. They find that when agents are allowed to join coalitions, they learn that grouping and behaving cooperatively with their coalition members results in higher benefits. In this paper we will also follow this approach.

*2.2.2. DCF over dynamic topologies.* Previous approaches focus on static networks. However, in most real-world situations, the topology of the network changes in response to the state of the network and the other way around, namely the state of the network changes in response to the topology. In fact, there is an increasing interest in games on adaptative networks and its influence in cooperation, where agents may improve

their topology (see [Gross and Blasius 2008] for a review), for example by changing the neighbors with whom they interact. Within these, the IPD has been broadly used as a test bench to study cooperation in complex networks. Zimmerman et al. [Zimmermann et al. 2004] present a model where agents imitate the strategy of their neighbors and where they are allowed to rewire their neighbors under certain fixed restrictions. Their model of cooperation with network plasticity leads to hierarchical topologies. Continuing their work, Eguiluz et al. [Eguiluz et al. 2005] focus on how possible dynamical mechanisms achieve a small world connectivity, also by allowing to imitate strategies and to change links to its neighbors. However, in their work agents imitate their best neighbors to decide their strategy, while in our approach agents learn how to behave based on their past interactions. The same happens with the neighbor selection, since they use fix strategies, while in our approach agents learn to which agent type it is better to rewire to.

Pacheco et al. [Pacheco et al. 2006] provide a model which incorporates decisions of individuals when establishing new links or giving up existing ones, by evaluating the productivity of their links. In our work each agent does not evaluate the productivity of each of the links. Instead, each agent decides to which agent type it provides a better payoff to link to. Fu et al. [Fu et al. 2009] propose a model of coevolutionary prisoner's dilemma that allows agents to either adjust their strategies or switch their defective partners. Thus, they show that partner switching is effective to stabilize cooperation. Although in a different realm (the investigation of tag-based coordination), Griffiths et al. [Griffiths and Luck 2010] show that partner switching (rewiring in their terms) can help increase coordination resilience in the face of malicious behavior. Along this line, in this paper using individual decisions on how to add or remove/replace a link is used to improve cooperation and avoid defective behavior.

*2.2.3. Coalitions with resources.* Apart from how coalitions are formed, how the members of a coalition establish cooperation is an important issue. This is even more important when modelling nowadays interconnected world, where agents own resources that they can trade. Indeed, social and economical trading agreements or cooperation play an important role in our society. That is why we are interested in the scenario where agents cannot only cooperate or defect, like in the classical IPD, but where they can also own tradable resources. From an economical point of view, Yee [Yee 2003] develops an evolutionary game model of property ownership and trade. Based on an evolutionarily model of animal territoriality, the author models human property ownership theoretically, showing that trading is evolutionary preferred over permanent ownership without trade.

Based on that work, Burguillo et al. [Burguillo and Peleteiro 2010] perform a study of the extended IPD-Possessor-Trader model. They present how evolutionary forces allow the emergence of different types of strategies using a spatial scenario. However, their work studies a grid and static scenario, which may not represent real world dynamic and complex topologies among agents and where no coalitions can be formed. Moreover, in such model, agents play independently, i.e., they cannot join coalitions to improve cooperation.

Finally, although the coalition-based mechanism in this paper is based on Axelrod approaches, there is also the idea of self governing institution due to Nobel Laureate Elinor Ostrom [Ostrom 1990]. In this line, Pitt et al. [Pitt et al. 2012a] model resource allocation in a network applying a formal characterization of Ostrom's socio-economic principles. Even if their focus and application domain is different, we can see that there is a correlation between their work and ours. We can regard Yee trading model as a linear public good game; coalition formation as cluster formation (Ostrom's principle 1: membership); deciding inner and outer behaviors as self-organization of cluster

rules; and rewiring as cluster movement. Therefore, it may be interesting to study similarities between the two models, to see how our results can contribute their problem, and vice-versa. However, this is out of the scope of this paper, but can be regarded as promising future work.

### 3. MODEL DESCRIPTION

In our model, we consider an agent population using a network as its interaction topology, where we model agents as nodes, and relations among them as edges. These agents interact with the peers in their social neighborhood, i.e., the agents to which they are linked, playing the Possessors-Traders game (agents are possessors or traders, see Sect. 3.1). Thus agents not only cooperate or defect, but they also have resources, with which they can trade. Moreover, to increase the cooperation level of the multi-agent system, agents can form coalitions, since group decisions (social) can result in a mutually beneficial cooperation that holds over time. Finally, agents do not have static neighborhoods, but they can change partners using rewiring. Thus during the game, additionally of trading resources, each agent must decide:

- To belong to a coalition or to be independent: Agents must decide whether being independent or in a coalition provides more benefits.
- To whom to rewire: As agents can change their neighbors, i.e., they rewire to improve their neighborhood, they have to decide to which agent to rewire.

Additionally, agents in a coalition act as a unity, i.e., all the agents of a coalition must decide which action to perform with agents belonging to their coalition (insiders) and with agents not belonging to it (outsiders). Thus, which is the coalition behavior, and how it is decided, is an important factor in the dynamics of the coalition.

In the following subsections we explain in detail the main agent decisions (see Algorithm 1): the trading strategies, the coalitions strategies, and the rewiring strategies.

---

#### ALGORITHM 1: Agent.Cycle

---

```

1: Payoffs = TradeAgainstAllneighbors()
2: Rewiring(Payoffs)
3: ReviseCoalition(Payoffs)

```

---

### 3.1. Trading strategies

In this section we describe the strategies that agents use to trade among them. It is primarily based on the model of property ownership and trade [Yee 2003; Burguillo and Peleteiro 2010]. This model is an extension of the Iterated Prisoner's Dilemma (IPD), i.e., trading is modeled as an extension of the IPD's cooperate and defect actions. In the model of property ownership and trade, there are two types of players: Possessors (P), which own a resource; and Traders (T), which sell and buy resources. In the following subsections we explain in detail the trading strategies of our game, that we define as Iterated Possessor-Trader (IPT).

*3.1.1. Ownership (Possessors).* A Possessor (P) is an agent owning a resource. Its strategy models the practice of ownership, i.e., the agent owns a resource with which it does not trade. Their behavior depends on whether they own a resource or they do not own it, i.e., if a possessor owns a resource, it acts as a defector, but if it does not, it cooperates (see [Yee 2003] for details). We show the Possessor (P) strategy in Algorithm 2.

---

**ALGORITHM 2:** Possesor Strategy

---

```

1: if owning(resource) then
2:   Defect
3: else
4:   Cooperate
5: end if

```

---

3.1.2. *Trading (Traders)*. A trader (T) is an agent willing to sell or buy a resource when dealing with a fellow trader. If an agent has the capability of trading, it uses it to try to maximize its benefits by selling a resource when owned. Indeed, nowadays society is based on trading agreements, obtaining services on the one hand, and benefits on the other. In particular, when two traders meet, the owner (agent that owns the resource to trade) values the resource at a random value  $y^2$ , where  $v < y < V$ ,  $(v, V) \in \mathbb{R}$ . Then, the buyer (the agent that wants to get the resource) offers a value  $x$  for the resource, where  $v < x < V$ . If  $x > y$ , then the buyer purchases the resource at random value  $z$ , being  $y < z \leq x$ . In [Yee 2003], the author models this norm introducing the Trader (T) strategy (see Algorithm 3). Note that if a trader plays against an agent that is not a trader, it behaves as a possessor (agents can know the type of the adversary before playing).

---

**ALGORITHM 3:** Trader Strategy

---

```

1: if is.Trader(neighbor) then
2:   if onwing(resource) AND  $v < y < V$  then
3:     Sell for y
4:   else if not(owning(resource)) AND  $y < x$  then
5:     Buy for z
6:   else
7:     Behave as Possessor
8:   end if
9: else
10:  Behave as Possessor
11: end if

```

---

3.1.3. *The Prisoner's Dilemma Game*. The model of property ownership and trade is an extension of the Iterated Prisoner's Dilemma (IPD). IPD models a situation in which two agents have to decide whether to cooperate (C) or defect (D), but without knowing what the other is going to do. In the IPD, the payoffs achieved in interaction are the following: if both agents cooperate, they get a reward ( $Re$ ) each, but if they both defect they get a punishment ( $Pu$ ). If one defects and the other cooperates, the first one gets  $Te$  (meaning the temptation payoff), and the cooperator receives  $Su$  (the sucker's payoff) (Table I). A prisoner's dilemma game satisfies the inequalities  $Te > Re > Pu > Su$  and  $2Re > Su + Te$ . Considering that  $Te > Re$ , it pays to defect if the other player cooperates. When the other player chooses defection, there is a choice between defection, that provides P (the punishment for mutual defection) or cooperation which yields  $Su$  (the sucker's payoff). Again, considering  $Pu > Su$ , it pays to defect if the other player defects. Thus, independently from what the other player does, it pays to defect. However, when they both defect, they get  $Pu$  instead of  $Re$ , which is a higher value that they both could get if they cooperate [Axelrod 1984].

<sup>2</sup>We choose a random value because the valuation of the resources is beyond the scope of this paper.

	Player $A_j$ Cooperates	Player $A_j$ Defects
Player $A_i$ Cooperates	Re, Re	Su, Te
Player $A_i$ Defects	Te, Su	Pu, Pu

Table I: General Prisoner’s Dilemma Matrix

Payoffs in this model are set based on [Yee 2003]. The author defines a payoff matrix for its trader game starting from the Hawk-Dove (HD) problem [Maynard Smith and Price 1973]. If  $\frac{1}{2}V > h$  then the HD corresponds to the PD game. As this is our case, we model our defection and cooperation based on the HD values that Yee defines for his Trader game:

$$\begin{aligned}
 Te &= (V + v)/2 \\
 Re &= (V + v)/4 \\
 Pu &= (V + v)/4 - h \\
 Su &= 0
 \end{aligned}$$

### 3.2. Coalitions

The basic strategy by which agents join a coalition or change to a new one is shown in Algorithm 4. If an agent ( $a_i$ ) has the worst payoff within its neighborhood after the last round (line 1), it joins the agent that has had the best one (line 2). If  $a_j$  is an independent agent, then  $a_i$  joins  $a_j$  to create a new coalition (line 4); but if  $a_j$  already belongs to a coalition, then  $a_i$  joins  $a_j$ ’s coalition (line 6). Note that this rule also enables any agent to change from a coalition to another in case it receives very poor payoffs in the former one.

In this dynamic network, agents form coalitions to act as a unity. All agents belonging to a coalition (each agent can belong to only one coalition at a time) do not need to be linked among them: they are a set of agents that act together, to maximize their performance. However, even though an agent does not have to have a link to all their coalition mates, it must have at least one link with one agent belonging to its coalition; if not, it becomes independent (lines 9-10, Algorithm 4). This is because if the agent is not connected to any of their coalition mates, it cannot know coalition information, strategy, share and divide gains, thus it must become independent. Again, notice that if agents change links, it does not imply that they change coalitions: rewiring to others means changing neighbors. We explain in following section how agents form and join coalitions, and how they decide the coalition behavior.

When agents are in a coalition, they must agree on the behavior to play with the other agents in the coalition (*insiders*) and with agents outside the coalition (*outsiders*). In our approach, we consider that coalitions are flat, i.e., there are no leaders, nor central authority imposing any policy, unlike in [Salazar et al. 2011; Burguillo-Rial 2009]. To decide the coalition behavior, each agent votes for a strategy (P or T) to play with *insiders* and for a strategy to play with *outsiders* (line 16, Algorithm 4).

To decide its vote, each agent uses a Learning Automata (LA) algorithm [Narendra, K.S. and Thathachar, M. A. L. 1989] that is trained from its trading history experience and payoffs. The LA algorithm keeps two probability models, one to assess the strategy to play against insiders ( $[ProbInT, ProbInP]$ ) and another to assess the strategy to play against outsiders ( $[ProbOutT, ProbOutP]$ ). Specifically, each agent uses Eq. 1 to reinforce the action with which the agent has obtained a higher payoff in the past:

**ALGORITHM 4:** ReviseCoalition(Payoffs)

---

```

1: if WorstPayoffInneighborhood(Payoffs) then
2:    $a_j = \text{neighborWithBestPayoff}()$ 
3:   if Independent( $a_j$ ) then
4:     CreateNewCoalition( $a_j, a_i$ )
5:   else
6:     JoinCoalition( $a_j, a_i$ )
7:   end if
8: end if
9: if Isolated() then
10:  GetIndependence()
11: else
12:  [ProbInT, ProbInP] = UpdateLAInsiders(Payoffs)
13:  [ProbOutT, ProbOutP] = UpdateLAOutsiders(Payoffs)
14:  ActionIn = ChooseInAction(ProbInT, ProbInP)
15:  ActionOut = ChooseOutAction(ProbOutT, ProbOutP)
16:  VoteBestBehavior(ActionIn, ActionOut)
17: end if

```

---

$$\forall_{j \neq i} : \begin{cases} p^{i,t+1} &= p^{i,t} + \alpha(1 - p^{i,t}) \\ p^{j,t+1} &= p^{j,t}(1 - \alpha) \end{cases} \quad (1)$$

In these equations,  $p^{i,t+1}$  is the probability that an agent performs a concrete action  $i$ , and  $\alpha \in [0, 1]$  is a (small) learning factor. The first rule is used to reinforce the action  $i$  chosen if it performed better than its alternatives in the considered state. At the same time, we apply the second rule to the other actions, decreasing its probability. In the next round, the agent chooses its new strategy using the updated probabilities. For instance, if the agent has obtained more gains in the past by playing action T against insiders (members of the coalition), then the probability of behaving as T with insiders is higher than behaving as P, and the same for playing against outsiders (out of the coalition). Using the corresponding probability model, each agent decides the action it wants to propose (line 14-15, *ActionIn* and *ActionOut*).

Once each agent has calculated its strategy (*ActionIn* and *ActionOut*), all the members of a coalition vote to decide the coalition strategy with insiders and outsiders (line 16, Algorithm 4). The voting is carried out using a voting protocol in which agents pass the vote, so there is no need of a central entity. Moreover, contrary to the case of having a leader, where agents belonging to a coalition may pay taxes to the leader [Burguillos-Rial 2009; Salazar et al. 2011], all the payoff within the coalition is equally and fairly shared among all coalition members. We choose a simple approach for the payoff distribution, since it is out of the scope of this paper to study other more complex approaches. However, there are several fair division techniques in social choice that are worth exploring, as in [Chevaleyre et al. 2007; Pitt et al. 2012b; Brandt et al. 2013]. We consider this as a promising line for future work.

### 3.3. Rewiring mechanism

In most real-world network interactions, relationships are not static, i.e., agents can change the individuals with whom they interact. This capability is modeled introducing a rewiring mechanism. Introducing rewiring, agents can modify their neighborhood whenever they are not satisfied with the outcome they receive from their actual neighbors. In the current work, we consider that only one link can be rewired by each agent at each iteration, although each agent can have several neighbors. We do not

**ALGORITHM 5:** Rewiring(Payoffs)

---

```

1:  $A_w = \text{WorstNeighbour}()$ 
2: if Independent() then
3:    $[\text{ProbRewP}, \text{ProbRewT}] = \text{UpdateLARewiring}(\text{Payoffs})$ 
4:    $A_c = \text{FindNewAgentToRewire}(\text{ProbRewP}, \text{ProbRewT})$ 
5: else
6:    $A_c = \text{BestNeighbourWithinCoalition}()$ 
7: end if
8: if ShouldIRewire() then
9:   if  $A_c \neq \emptyset$  AND AcceptsRewiring( $A_c$ ) then
10:     $\text{LeaveWorstNeighbour}(A_w)$ 
11:     $\text{RewireTo}(A_c)$ 
12:   end if
13: end if

```

---

consider the all-or-nothing rewiring, where an agent adopts a new neighborhood, since that extreme rewiring is not realistic in most scenarios [Griffiths and Luck 2010].

We have two rewiring strategies, depending on whether the agent is independent or belongs to a coalition:

**(1) Independent agents**

Independent agents use a Learning Automata to learn which type of agent it is the best to rewire and the kind of agents must be refused for rewiring. Analogously to the coalition strategy, each agent keeps track of its history and revises the decision depending on its experience and payoffs. The LA stores a probability model  $[\text{ProbRewP}, \text{ProbRewT}]$  which represents the probability to rewire to a possessor or to a trader. Specifically, the probability model is updated (line 3, Algorithm 5) using Eq.1, reinforcing to rewire to the type of agents that has provided more gain in the past. Once every independent agent has updated the probability vector, it looks for a candidate agent to rewire fitting the current probability model ( $A_c$ , line 4, Algorithm 5). Since there are two types of agents (P or T), the selected candidate is randomly selected within all the population of the estimated best type (P or T). However, it can happen that there is not a candidate, or the selected agent  $A_c$  refuses the agent that wants to add a link with it, if it does not belong to the type of agents that they learn that they want to rewire. Finally, note that since the number of links in the network must remain constant (we do it as in [Griffiths and Luck 2010]), if the rewiring succeeds, the agent must leave the neighbor with the worst payoff in the last round (line 10, Algorithm 5).

**(2) Coalition agents**

A coalition agent may rewire to the neighbor's best neighbor. This greedy approach takes advantage of the qualified information that coalition agents may obtain from their peers. Moreover, by allowing to rewire to agents that may not belong to the coalition, we provide more flexibility for coalition formation and agent interaction. To rewire, a coalition agent asks all its coalition peers for the neighbor that provided the biggest payoff in the last round (line 6, Algorithm 5). Then, the agent rewires to the best neighbors' neighbor ( $A_c$ ) and leaves the peer that provided it the worst payoff ( $A_w$ ).

**4. EXPERIMENTS**

The aim of this section is to empirically analyze the cooperation mechanism for the resource trading environment introduced in Section 3. First experiments focus on analysing the contribution of the coalition and rewiring mechanisms to increase the

payoffs of the agents. Specifically, each proposed mechanism is compared with the complete coalition+rewiring proposal. Moreover, we compare our approach against the mechanism described in [Salazar et al. 2011]. We will also focus our attention on the dynamics of coalition formation and on the influence of rewiring in the coalition formation process. To evaluate the resilience of our approach, we will also report experiments on the presence of free-riders (agents that, despite belonging to a coalition, do not follow its agreements).

Hence, first in Section 4.2 we start by comparing four cooperation mechanisms, namely:

- *Base*. Agents employ the Iterated Possessor-Trader (IPT) as their interaction mechanism.
- *Rewiring-only*. Agents interact by means of the IPT, but they also use rewiring to change their neighborhoods.
- *Coalition-only*. Agents interact by means of IPT, but they also use coalition decisions to either form, join or leave coalitions.
- *Coalition-plus-rewiring*. Agents interact by means of IPT, but they also use coalition decisions and rewiring.

Our purpose is to show the benefits in terms of payoff increase that agents derive when using the coalition plus rewiring mechanism with respect to the other three cooperation mechanisms.

Next, in Section 4.3 we analyze the dynamics of the coalition formation mechanism described in Section 3 when paired with rewiring. We study how the coalition formation mechanism partitions the agent population into coalition agents (those belonging to some coalition) and independent agents. Our purpose is to quantify the effectiveness of our coalition formation mechanism in helping agents to form coalitions. Furthermore, we also analyze the effect of adding rewiring to our coalition formation cooperation mechanism. More precisely, we observe how rewiring influences the coalition formation process by observing the number and sizes of the coalitions formed by agents.

Then, in Section 4.4 we focus on the coalition-plus-rewiring mechanism identified as the one leading to higher gains by our first experiment. Then, we explain the evolutionary behavior of agents within coalitions: (i) to understand how they strategically behave with other agents inside and outside their coalitions; and (ii) to understand their rewiring behavior. The purpose of this analysis is to observe how agents learn to be more cooperative.

In Sect. 4.5 we analyze the effect of varying our payoff matrix while maintaining the constraints that payoff values must satisfy. Our purpose is to show that our results do not depend on the specific parameters we have chosen for the represented experiments.

Finally, in Sect. 4.6, we study the influence of rewiring in the coalition formation process. We do this since rewiring and coalition formation are closely related, and, as we will show in our experiments, rewiring has a non-trivial effect on coalition formation.

#### 4.1. Empirical settings

Our empirical evaluation is based on a discrete-event simulation of a population of agents interacting with one another on a network. Each agent is placed in a node of the network and two agents cannot be placed at the very same node. Interactions only occur between a pair of agents whenever they are connected by a link of the network.

At the outset of a simulation, agents in the agent population are endowed with one of the cooperation mechanisms listed at the beginning of Section 4. A simulation consists of a sequence of simulation steps. At each simulation step, each agent will be able to interact with its neighboring agents playing the game described in Section 3 (either

as part of a coalition or as an independent agent). Furthermore, depending on the cooperation mechanism deployed on the agent population, each agent will be allowed to form coalitions, join coalitions, leave coalitions, and rewire to other agents. As a simulation proceeds, each agent accumulates payoffs as a result of the games played with its neighboring agents. We shall consider that a simulation has converged when the number of coalitions formed by the agent population remains stable after twenty simulation steps.

For each of the experiments reported below, we ran ten simulations till convergence. The reported results correspond to averaging the values observed during these simulations unless indicated otherwise. Moreover, we show the variances of average values.

*4.1.1. Agent population.* Unless stated otherwise, for each experiment we generate an agent population composed of 400 agents. The IPD game matrix employed by all agents in the population is set as follows:  $Te = 2.5$ ,  $Re = 1.25$ ,  $Pu = 0.25$ , and  $Su = 0$ . Each agent's initial strategy to play in a game is chosen at random so that there is a 0.5 probability that each agent is initially either a trader (T) or a possessor (P). Moreover, an agent's initial strategy is the one played by its when it becomes independent. The trading values required by traders ( $v$  and  $V$ ) are sampled from a uniform distribution  $U[0, 5]$ . We set  $h = 1$ .

Recall from Section 3 that possessors and traders behave differently depending on whether they own resources or not. Thus we will generate agent populations with varying distributions of resources: from 10% up to 90% of the agents in a population owning resources. We will be particularly interested in investigating three types of cooperation scenarios, which depend on the following distributions of resources in the population:

- *Scarcity of resources.* A small fraction of the agent population (10%) own resources.
- *Balanced resources.* Half of the agents in an agent population (50%) own resources.
- *Plentiful resources.* There is plenty of agents owning resources in the agent population (90%).

Finally, the learning factors used by all agents employing coalition strategies and rewiring strategies,  $\alpha_{coalition}$  and  $\alpha_{rewire}$  respectively, are both set to 0.1.

*4.1.2. Network topology.* We have chosen two types of network topologies: small-world and scale-free. Both small-world and scale-free networks provide realistic models of the topological features found in many nature, social, and technological networks [Reka and Barabási 2002; Pastor-Satorras and Vespignani 2001; Watts and Strogatz 1998].

On the one hand, small-world networks model real world complex systems such as neural networks, food webs, scientific-collaboration networks, and computer networks [Kniesburges et al. 2012]. These networks are characterized by the small-world phenomenon, in which nodes have small neighborhoods, and yet it is possible to reach any other node in a small number of hops. This type of networks is highly-clustered, namely they have a high clustering coefficient. Recall that the clustering coefficient is a measure of degree to which nodes in a graph tend to cluster together. Thus, small-world networks tend to contain cliques, and near-cliques, meaning sub-networks that have connections between almost any pair of nodes within them. Formally, we note a small-world network as  $W_N^{k;p}$ , where  $N$  is the number of nodes,  $k$  stands for the average connectivity (the average size of a node's neighborhood), and  $p$  is the rewiring probability. In this paper we employed the Watts & Strogatz model [Watts and Strogatz 1998] to generate small-world networks with the following settings:  $N = 400$ ,  $k = 5$ , and  $p = 0.1$ .

On the other hand, scale-free networks model real-world networks like the world-wide web [Adamic, L.A. and Huberman, B.A. 2000], the Internet, and some biological networks [Newman 2003]. These networks are characterized by having a few nodes acting as highly-connected hubs, while the rest of them have a low connectivity degree. Unlike small-world networks, scale-free networks are low-clustered networks. Formally we note a scale-free network as  $S_N^{k;-\gamma}$ , where  $N$  is the number of nodes, and the probability that a node in the network connects with  $k$  other nodes is roughly proportional to  $k^{-\gamma}$ , namely  $P(k) \sim k^{-\gamma}$ . In this paper, we employed the Barabasi-Albert algorithm [Reka and Barabási 2002] to generate scale-free networks with the following settings:  $N = 400$ ,  $k = 5$ , and  $\gamma = 2$ .

#### 4.2. Evaluating cooperation mechanisms

Next we compare the *coalition-plus-rewiring* mechanism with the *base*, *rewiring-only*, and *coalition-only* mechanisms. Our aim is to quantify the benefits in terms of payoff that agents obtain when employing the *coalition-plus-rewiring* mechanism. We measure such benefits depending on the distribution of resources in the agent populations that we simulate. More precisely, we assess the benefit of a given cooperation mechanism as the percentage of payoff that it obtains with respect to the maximum payoff that it could obtain. We obtain the maximum payoff from the simulation run with the maximum average payoff out of all the simulation runs performed for the four cooperation mechanisms.

Figure 1 shows the percentage of payoff gain obtained by the four cooperation mechanisms over different network topologies, and depending on availability of resources. Overall, all cooperation mechanisms lead to higher payoffs as more resources are available, namely the higher the availability of resources, the larger the profit obtained by cooperating. However, the slope of *rewiring-only*, *coalition-only*, and *coalition-plus-rewiring* is larger than the slope of *base*. This indicates that the increase of payoff with the availability of resources is larger when either using rewiring or coalitions or both. As we will discuss further in Section 4.4, as more resources are available, there are more opportunities to trade, and hence trading emerges as the preferred strategy. Moreover, independently of the topology, the *coalition-plus-rewiring* mechanism leads to the higher payoffs. However, the amount of profit that *coalition-plus-rewiring* delivers does differ depending on the network topology. As to small-world (see Figure 1a), observe that using *coalition-plus-rewiring* leads to up to 30% more payoff than the *base* mechanism over a small-world network (see Figure 1a) and up to 25% over a scale-free network (see Figure 1b).

So far we have compared our coalition-plus-rewiring mechanism against the simpler cooperation mechanisms described at the beginning of Section 4. Next, we compare our approach against the mechanism described in [Salazar et al. 2011]. In [Salazar et al. 2011] Salazar et al. introduce a dynamic coalition formation mechanism aimed at emerging cooperation. Unlike our approach, the mechanism in [Salazar et al. 2011] relies on leaders that impose the behavior of the coalitions and charge taxes to coalition members. Fig. 2 shows the percentage of gain of our mechanism (labelled *Without leaders* in the figure) against the leader-based mechanism in [Salazar et al. 2011] (labelled *With leaders* in the figure) as resource availability increases. We observe that our mechanism outperforms the leader-based mechanism. The larger the resource availability, the larger the percentage of gains of our mechanism with respect to the leader-based one. There are two main reasons that explain why our mechanism outperforms the leader-based one. First, agents using our mechanism do not pay taxes to any leader, thus all gains are shared among coalition members. Second, recall that leaders impose coalitions' behaviors. In doing so, a leader disregards valuable, local in-

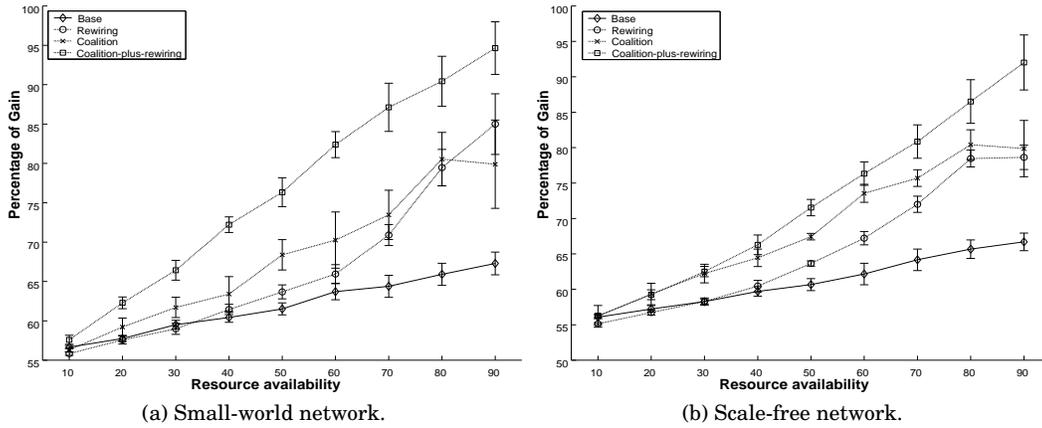


Fig. 1: Comparison of gains obtained by all cooperation mechanisms (with  $p_{rew} = 0.4$ ).

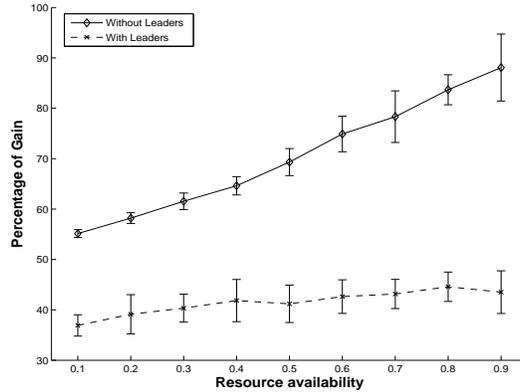


Fig. 2: Comparison of our mechanism with other previous approach that uses leaders (scale-free network and  $p_{rew} = 0.4$ ).

formation that might potentially benefit all coalition members, not only the leader. Unlike the leader-based approach, our mechanism lets agents learn their best coalition strategies from their own experience. Hence, agents can exploit such information to reach a consensus and decide the overall behavior of their coalitions.

Next we focus on the *coalition-plus-rewiring* mechanism to assess the effect of using rewiring on the payoff gain obtained. Figure 3 shows the percentage of payoff gain obtained depending on the rewiring probability employed by agents ( $p_{rew}$ ) and depending on the availability of resources (scarcity, balance, plentiful). Notice that the results when there is no rewiring ( $p_{rew} = 0$ ) correspond to using the *coalition-only* mechanism. Recall that the use of rewiring is intended to help an agent increase its payoffs by disconnecting from neighbors that have demonstrated not to deliver good payoffs in the past and connecting to other agents that may potentially deliver better payoffs in the future. We observe that indeed rewiring leads to increase gains independently of the topology, and that by introducing small rewiring ( $p_{rew} = 0.1$ ) the gains are al-

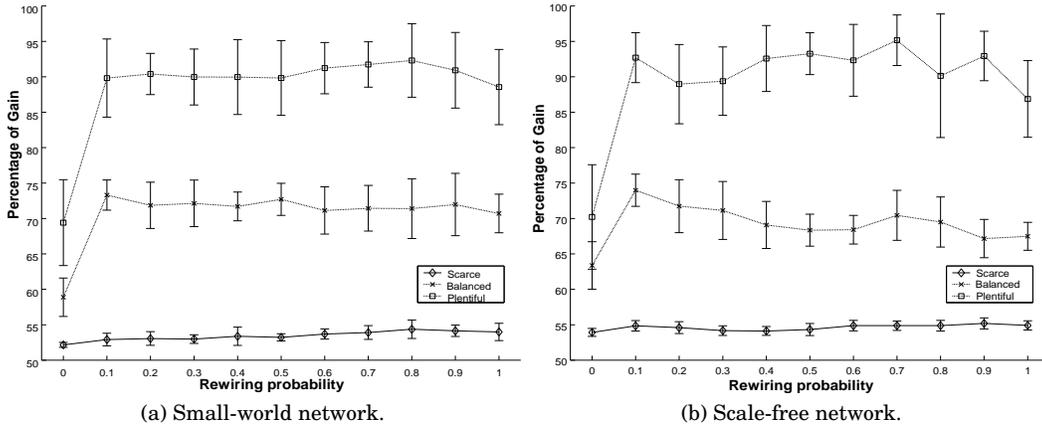


Fig. 3: Coalition-plus-rewiring mechanism. Percentage of payoff gain per agent when varying the rewiring probability of agents and depending on the availability of resources (scarcity, balance, plentiful).

ready highly increased (20% with plentiful resources). Moreover, this payoff gain does depend on the distribution of resources: the larger the number of resources, the larger the payoff gain. We observe for both topologies that total rewiring ( $p_{rew} = 1$ ) represents an extreme case depending on the distribution of resources. Thus, using total rewiring (continuously changing neighbors) has a rather negative effect when there is either plentiful or balanced resources. However, although total rewiring is worse in terms of payoff gain than lower rewiring probabilities, it is still better than no rewiring at all. Notice that agents benefit from total rewiring when there is low resources in both topologies. This is because the fact that agents continuously change neighbors help them to find agents eager to trade in a scenario where resources are scant.

Finally, we are also interested in studying the sensitivity of our mechanism to free riders (agents that, despite belonging to a coalition, do not follow its rules). Fig. 4 shows the percentage of gain obtained when using our mechanism without free riders compared to the percentage of gain obtained as the percentage of free riders increases. Observe that as the percentage of free riders increases, the percentage of gain decreases. The gain loss ranges from  $\sim 14\%$  to  $\sim 33\%$ . Thus, free riders do affect our mechanism. However, notice that the gain loss is not dramatic. This is because our mechanism manages to isolate a free rider within a coalition because its members realize that the free rider is selfishly looking for benefitting from the coalition. Thus, the free rider ends up not belonging to the coalition. However, since we assume anonymity (agents are not identified with ids), after an agent is expelled from a coalition, she is free to join another one because this new coalition knows nothing about the agent's past. Hence, future work should aim at extending our mechanism to cope with anonymous free-riders.

### 4.3. Analyzing coalition formation dynamics

As we have shown in the previous section, grouping agents into coalitions and rewiring improves their payoff gains. Hereafter we analyze how the coalition formation mechanism, described in Section 3, partitions the agent population into coalitions agents and independent agents. We pursue to measure the effectiveness of our coalition for-

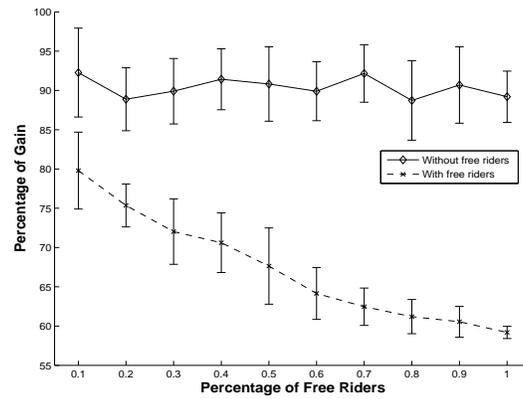


Fig. 4: Gains obtained by our mechanism when the agent population contains free riders and when it does not (scale-free network,  $p_{rew} = 0.4$  and plentiful resources (90%)).

mation mechanism in helping agents to form coalitions. Moreover, we analyze how rewiring influences the coalition formation process by observing the features present in the coalitions that emerge.

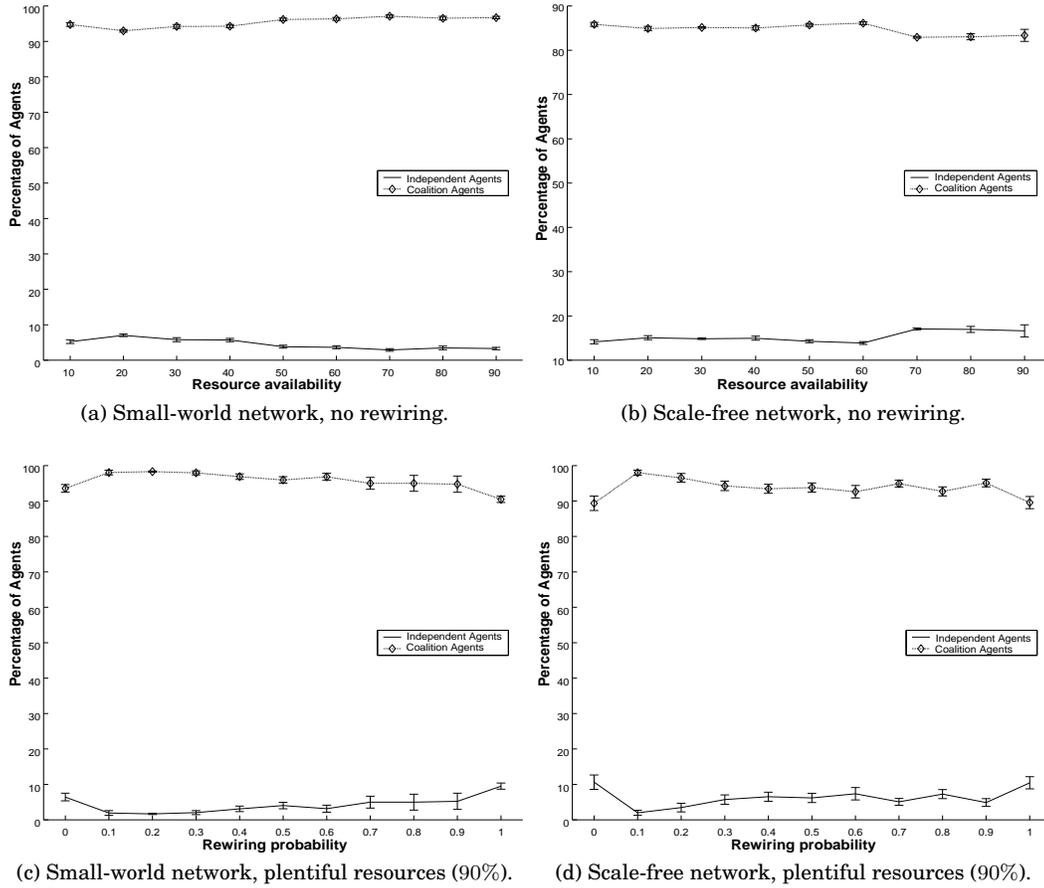


Fig. 5: Independent versus coalition agents.

According to our model in Section 3, agents can decide either to stay independent or join a coalition. Figure 5a and Figure 5b show the percentage of agents belonging to coalitions depending on the distribution of resources on small-world and scale-free topologies. Since agents employ no rewiring, both cases correspond to using the *coalition-only* cooperation mechanism. Observe that the results are slightly affected by the amount of resources in the agent population. On the one hand, around 90% of the agents in the population become coalition agents on small-world networks, whereas around 85% of the agents become coalition agents on scale-free networks. Figure 5c and Figure 5d show the percentage of coalition agents when using rewiring. For both topologies, the use of rewiring leads to a 5 – 10% increase in the percentage of coalition agents. Thus, the percentage of coalition agents comes close to 95% on small-world networks and close to 95 – 100% on scale-free networks. Notice though that the percentage of coalition agents decreases with total rewiring ( $p_{rew} = 1$ ). Thus, total rewiring has a detrimental effect on the coalition formation process, since when agents constantly change their neighbors, rewiring acts as noise. Moreover, we have observed, that these individual agents are ostracized. Since they do not show a cooperative behavior, other

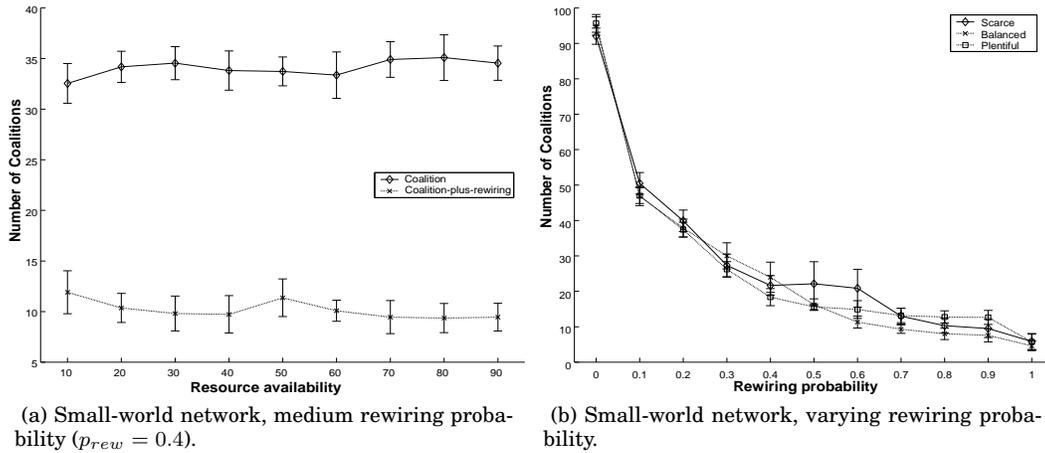


Fig. 6: Number of coalitions.

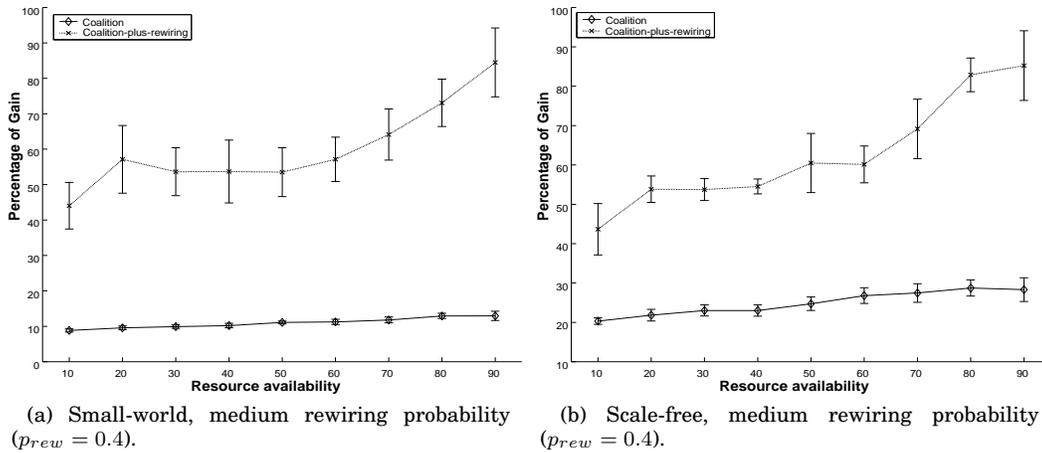


Fig. 7: Average payoff gain per coalition.

agents rewire to other agents, isolating them. This is in line with the results of Fu et al. [Fu et al. 2009], where they show that isolated individuals are often defectors.

Now we turn our attention to analyzing how our coalition formation mechanism partitions the agent population into different coalitions. Figure 6a shows the number of coalitions as the percentage of resources in the population varies for the *coalition-only* and *coalition-plus-rewiring* cooperation mechanisms. Observe that allowing agents to rewire leads to significantly reduce the number of coalitions that are created by means of *coalition-only*. In fact, *coalition-plus-rewiring* leads to less than a half of the coalitions created by means of *coalition-only*. Therefore, using rewiring has a compacting effect on the coalition formation process: less and bigger coalitions. Since agents within the same coalition cooperate, having less and bigger coalitions is bound to yield significant payoff benefits. This is confirmed by Figure 7, which shows the percentage of pay-

off gain that agents obtain when employing *coalition-plus-rewiring* versus employing *coalition-only*. We observe that the benefits are larger as the percentage of resources in the agent population is higher, thus our approach takes advantage of the availability of resources to increase the benefits of the agents.

In general, as shown by Figure 6b, the larger the rewiring probability, the lower the number of coalitions. Notice also that when there are abundant resources (90%), the number of coalitions reaches its minimum. This occurs because, in general, the larger the number of resources, the larger the number of possessors behaving as defectors, and hence agents learn to group into bigger coalitions to avoid defection.

So far we have shown that the number of coalitions decreases as the rewiring probability employed by agents increases. Hereafter we show that the size of the coalitions formed increases with rewiring. Figure 8 displays several histograms representing the number of coalitions that are formed per coalition size as the rewiring probability increases for both scale-free and small-world networks. Figures 8a, 8b, and 8c show how the size of coalitions formed by agents increases as rewiring increases. Thus, in Figure 8a, there are 45 coalitions with 2 or less than 2 agents, being 22 the maximum coalition size. In Figure 8b there are 15 coalitions with less than 10 agents, and the maximum coalition size goes up to 80. Finally, in Figure 8c, a big coalition composed of more than 250 agents is formed, while the rest of agents spread into coalitions with less than 50 agents (e.g., there are 92-agent coalitions). Rewiring leads to this increase in coalition sizes. Since agents are allowed to change their neighbors, they can choose to wire to agents that provide more benefits and to join bigger coalitions to be more efficient against non-cooperative behaviors. Together, the decrease in the number of coalitions, and the increase in their size, results in higher payoffs for agents. Notice that we observe the same effect of rewiring on small-world networks, as shown by Figures 8d, 8e, and 8f.

#### 4.4. Analysing agents' behaviors

Now we focus on the individual behavior of agents employing the *coalition-plus-rewiring* mechanism. We dissect the individual behavior of agents: (i) to understand how they strategically behave with other agents inside and outside coalitions; and (ii) to understand their rewiring behavior. Our aim is to learn how agents' individual behaviors lead to cooperation.

*4.4.1. Coalition behavior.* Recall from Section 3.2 that agents within a coalition vote to decide the coalition behavior to play both against insiders (coalition members) and outsiders (agents out of the coalition). Such decisions depend on what agents learnt in the past, i.e., on the actions that yield more payoffs. Note that agents choose the action with the highest number of votes. In case of draw, they randomly choose one of the tied actions. Therefore, consensus is always guaranteed. Next we study which actions agents choose to improve their payoffs.

Figures 9a and 9b show the percentage of coalition agents per strategy (possessor or trader) for both small-world and scale-free topologies as the distribution of resources varies. As resource availability increases, coalition agents' behaviors both against outsiders and insiders (coalition mates) varies. Thus, along the line of [Fu et al. 2012], we are interested in the conditions under which preferential in-coalition cooperation emerges. Moreover, we are also concerned with assessing when behaving cooperatively with outsiders of a coalition might be beneficial.

We consider first the small-world topology. In Figure 9a we clearly differentiate three scenarios:

— *Low cooperation.* When there is less than 20% of resources, coalition agents do not trade with outsiders, but mostly do it with insiders (around 75% of the agents behave

as traders with coalition-mates). Since there are few resources to trade with, even behaving as possessors with coalition-mates eventually occurs. Nonetheless, coalition agents quickly learn that it is more beneficial to trade with coalition mates. To summarize, the dominant strategy with outsiders when the percentage of resources is very low is to be a possessor.

- *No dominant strategy.* When the percentage of resources is around 20%, there is a transition in the behavior of coalitions with outsiders. Thus, we observe that half of coalitions behave as possessors with outsiders, while the other half behave as traders. At this point there is no dominant strategy with outsiders.
- *High cooperation.* Beyond 20% of resources, coalition agents progressively become more and more cooperative with outsiders. A medium or large availability of resources allows agents to perform more trades, obtaining higher benefits by cooperating in this way. However, notice that the percentage of traders with insiders is larger than with outsiders, though the gap closes as the availability of resources increases because the number of trades increases and also the defective behavior of outsiders. To summarize, trading becomes the dominant strategy against outsiders.

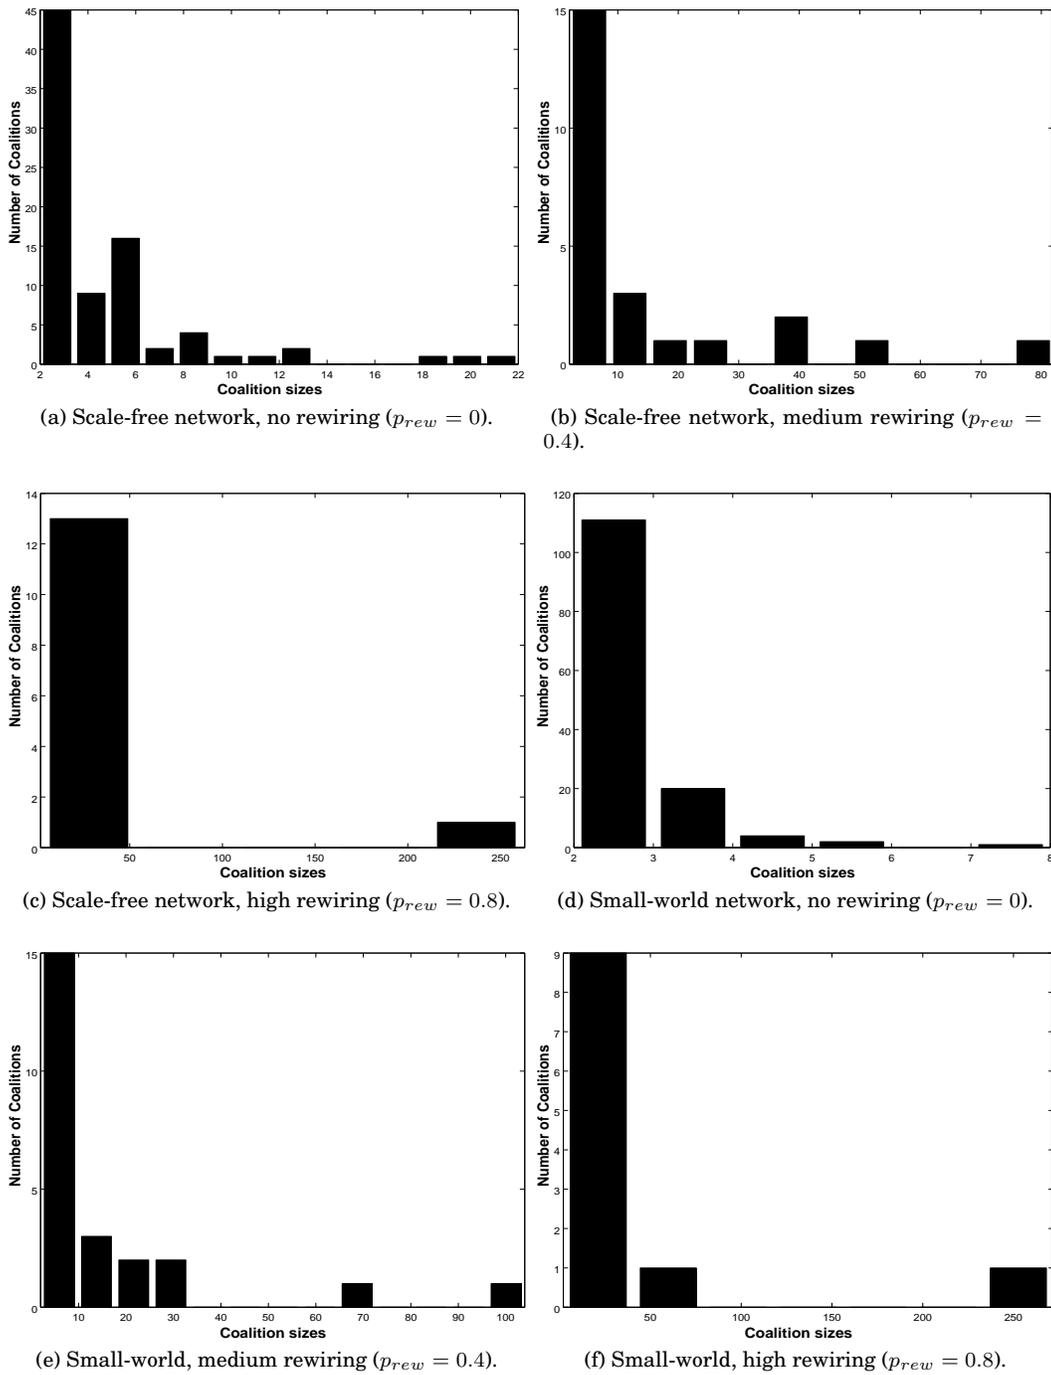


Fig. 8: Histogram for the size of coalitions created when employing the *coalition-plus-rewiring* cooperation mechanism. The x-axis represents the size of coalitions, and the y-axis represents the number of coalitions.

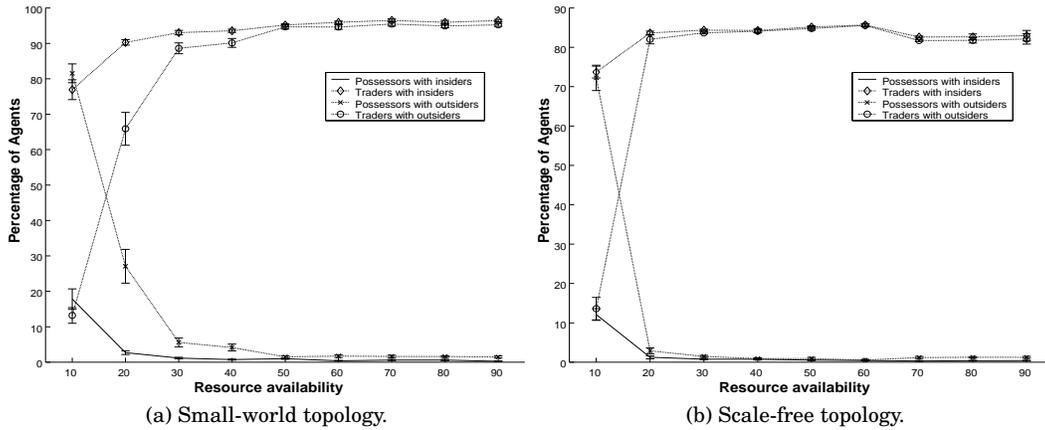


Fig. 9: Percentage of agents per strategy within coalitions.

Consider now the scale-free topology. We observe the very same scenarios described above. However, there are slight differences. Firstly, in the low cooperation scenario, the percentage of traders with outsiders is larger (hence the percentage of possessors, is smaller). Secondly, once the transition between low and high cooperation occurs (beyond 20% of resources), trading with outsiders very quickly becomes the dominant strategy.

**4.4.2. Rewiring behavior.** Recall that the rewiring strategy described in Section 3 allows agents to change their neighborhoods. Recall also from Section 4.2 that using rewiring together with coalition formation, the so-called *coalition-plus-rewiring* cooperation mechanism, helps agents to obtain higher payoffs than only using coalition formation (as shown in Figure 1).

First, we analyze the rewiring behavior of coalition agents. Figure 10 shows the distribution of links for coalition agents with coalition-mates and with outsiders for both small-world and scale-free topologies. The rewiring behavior radically differs depending on the network topology.

On the one hand, regarding small-world networks, we observe that as the rewiring probability increases, the number of links that agents establish with coalition-mates increases. Thus, when the probability of rewiring is large (beyond 80%), coalition agents establish more links with coalition-mates than with outsiders. This results in fewer and bigger coalitions, as shown in Figure 8, which are loosely connected with outsiders. Therefore, cooperation increases and agents obtain larger payoffs, as shown in Figure 3a<sup>3</sup>.

On the other hand, we do not observe the same behavior on scale-free networks. Indeed, as the rewiring probability increases, the number of links that agents establish with coalition-mates increases too. However, the percentage of links with outsiders remains significantly larger. To understand this differing behavior, we further investigated how the topology of a scale-free network evolves. Figure 11 illustrates the final topology of a network (depicted using Pajek [Batagelj and Mrvar 2003]) that started as a scale-free network upon convergence. The network contains 100 agents (to ease

<sup>3</sup>Except when agents employ total rewiring, which has a detrimental effect, as discussed in Section 4.2.

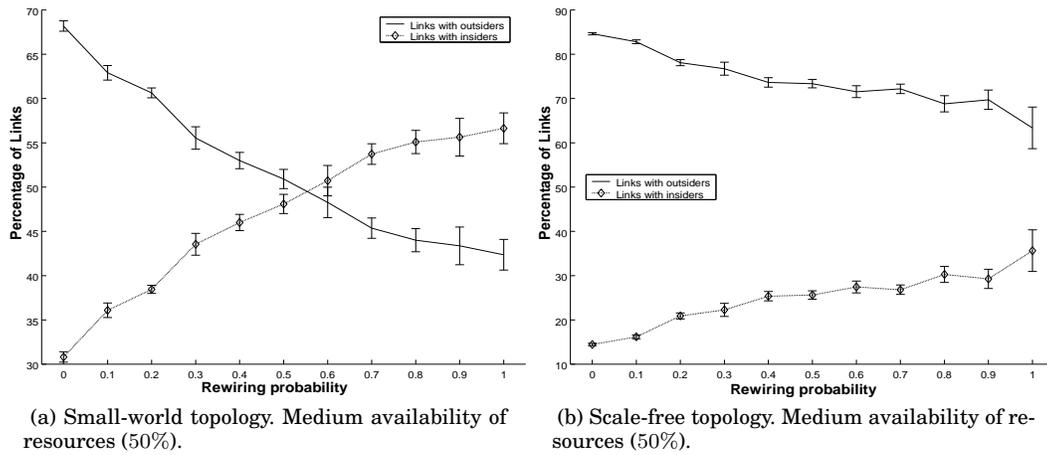


Fig. 10: Percentage of links of coalition agents: with insiders (coalition-mates) and with outsiders.

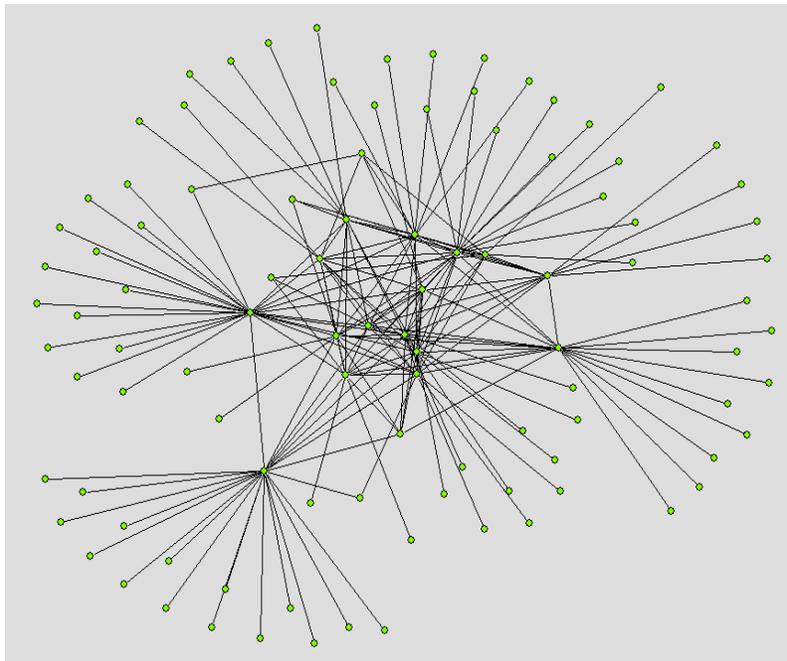


Fig. 11: Final topology of a scale-free network after agents deploy the *coalition-plus-rewiring* mechanism.

its display) with high availability of resources (90%) and medium rewiring probability ( $p_{rew} = 0.4$ ). Two different types of coalitions are formed:

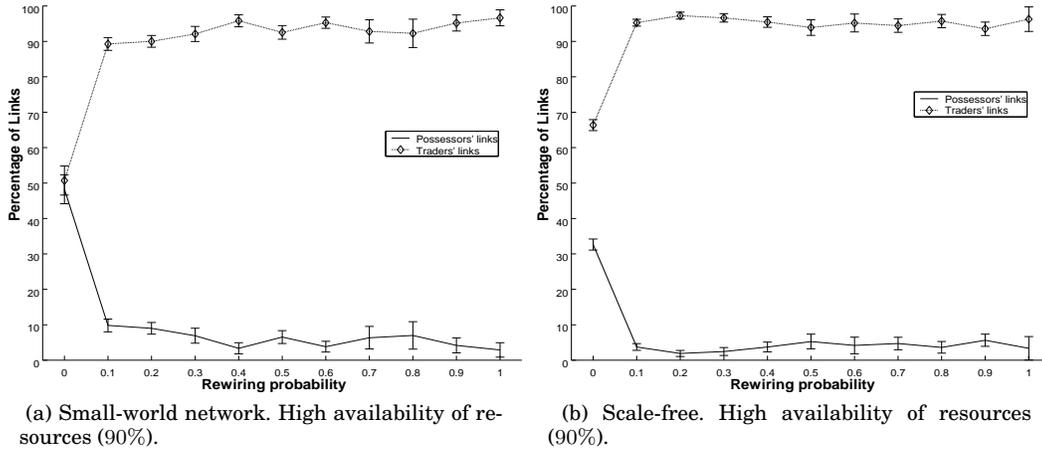


Fig. 12: Percentage of links of coalition agents with outsiders.

- *Hub-centered coalitions.* A coalition that is organized around a hub. This hub is connected to a large number of agents. Each of these agents has in turn a single link to the hub.
- *Clique-like coalitions.* These are coalitions whose agents share many links with coalition-mates, thus showing a clique-like structure.

We observe that the number of hub-centered coalitions is larger than the number of clique-like coalitions. Therefore, this explains why the percentage of links with insiders is low compared to the percentage of links with outsiders, as shown in Figure 10b.

Finally, we analyze how coalition agents establish links with outsiders. Figure 12 shows the percentage of links established by coalition agents with outsiders depending on the outsiders' cooperation behavior in an environment with plentiful resources, i.e., a highly defective environment in the presence of possessors. The figures show how the percentage of links varies with the rewiring probability for both small-world and scale-free topologies. In both networks, we observe that agents behaving as traders with outsiders increase their number of links. The reason for this is that agents want to have as neighbors the agents that provide the largest payoffs. With this aim, an individual agent learns that rewiring to traders provides larger benefits, while a coalition agent recommends to rewire to its neighbors' traders (since these provide its with the largest benefits). This increases the links of traders with outsiders which, in turn, reinforces their trading strategy with outsiders since their gains are expected to increase.

#### 4.5. Discussion on the effects of varying payoffs

In the above-described experiments we have employed a payoff matrix based on the one introduced in [Yee 2003]. In this section we analyze the effect of varying our payoff matrix while maintaining the constraints that payoff values must satisfy as stated in Sect.3.1. Overall, our empirical evaluation indicates that changing the payoff matrix does not have a big influence in the final results (provided PD inequalities are maintained).

First, we observed the effects of increasing the value of  $Pu$  until it becomes almost same as  $Re$ . In this case, although the number of agents behaving as possessors with

insiders increases a 10%, trading prevails as the dominant strategy. Secondly, when increasing the value of  $Re$  until it becomes almost the same as  $Te$ , there is an increase of 20% of traders when resources are scarce. This result is coherent: when two traders meet and they do not have a resource, they both cooperate, obtaining  $Re$  as payoff. Finally, we have also investigated the effects of changing the value  $V$ . With this aim, we doubled its initial value. This caused that the percentage of agents of each type remains almost the same as in Fig.9. Nonetheless, we also observed that the total gains of the population increase, being this increase higher when there are more resources available.

#### 4.6. Effects of rewiring on coalition formation

As mentioned above, rewiring and coalition formation are closely related. Thus, we next focus on analyzing the effect of rewiring on coalition formation. Fig. 13 illustrates the influence of rewiring on the the number of agents that change its coalition (*Coalition changes*) versus the number of agents that change a link (*Rewiring changes*) for different rewiring probabilities.

Fig.13a shows the percentage of coalition changes and of rewiring changes along three different simulations corresponding to different rewiring probabilities (low, medium, and high, namely  $p_{rew} \in \{0.2, 0.4, 0.6\}$ ). As the probability of rewiring increases, not only the number of partner changes (rewirings) increases, but also the number of coalition changes. This is because the increase of rewiring leads to higher coalition instability. Thus, although an agent can be stable in a coalition, after rewiring she may discover a better coalition or become isolated from the coalition she belonged to. Furthermore, we observe that the gap between the number of rewiring changes and coalition changes also increases with the rewiring probability. That is, as we increase the probability of rewiring, the number of partner changes proportionally increases. However, coalition instability is lower than rewiring changes.

Fig.13b shows the effect of varying the value of the probability of rewiring during a simulation. At the outset of the simulation, we set the probability of rewiring to a large value ( $p_{rew} = 0.8$ ). After 50 simulation ticks, we set the probability of rewiring to 1/3 of its initial value. Thereafter the number of rewiring and coalition changes drop until we turn up again the the probability of rewiring: we set it to 2/3 of its initial value when reaching 100 simulation ticks. We may observe that both coalition and rewiring changes increase again. After 150 simulation ticks we set back again the probability of rewiring to 1/3 of its initial value. Finally, after 200 simulation ticks, we set the probability of rewiring to its initial value. Experiments reported show that the number of coalition and rewiring changes increases again to reach similar percentages to the ones at the beginning of the simulation.

#### 4.7. Summary

In this section we summarize the main results of the presented experiments. Unless otherwise stated, the results are the ones when using our mechanism *coalition-plus-rewiring*.

- (1) Analysis of agents' gains.
  - *Comparison of coalition-plus-rewiring with alternative mechanisms (base, rewiring-only, coalition-only)*. Overall, there is an increase of benefits with increase of resources. Moreover, our mechanism *coalition-plus-rewiring* leads to higher payoffs than the others.
  - *Comparison of coalition-plus-rewiring with a leader-based mechanism*. Our mechanism outperforms the leader-based mechanism in terms of agents' gains.

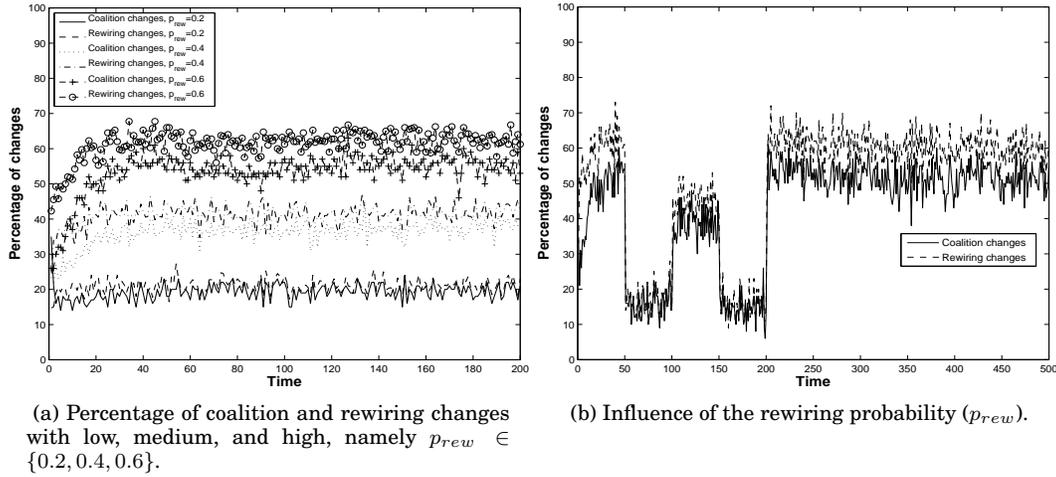


Fig. 13: Influence of rewiring on coalition formation

The larger the resource availability, the larger the percentage of gains of our mechanism with respect to the leader-based one.

- *Effects of rewiring.* Rewiring leads to increase gains independently of the topology when using *coalition-plus-rewiring*. In fact, introducing a small rewiring ( $p_{rew} = 0.1$ ) highly increases the gains already.
- *Sensitivity of coalition-plus-rewiring to free riders.* As the percentage of free riders increases, the percentage of gain decreases, ranging the loss between 14% and 33%.

(2) Analysis of the effects of rewiring on coalition formation.

- *Independent and coalition agents.* Using *coalition-only* (without rewiring), around 90% of agents are in a coalition. Using our mechanism *coalition-plus-rewiring* leads to a 5-10% increase in the percentage of coalition agents. Total rewiring ( $p_{rew} = 1$ ) has a detrimental effect on the coalition formation process because rewiring becomes noise when agents constantly change their neighbors.
- *Effects of rewiring over coalitions.* Using our mechanism *coalitions-plus-rewiring* leads to significantly reduce the number of coalitions that are created when compared to *coalitions-only*, in fact to less than half. Thus, the larger the rewiring probability, the lower the number of coalitions. Moreover, rewiring also leads to an increase in coalition sizes. Together, the decrease in the number of coalitions, and the increase in their size, results in higher payoffs for agents. Therefore, using rewiring has a compacting effect on the coalition formation process: less and bigger coalitions.
- *Effects of rewiring on coalition formation.* As the probability of rewiring increases, not only the number of partner changes (rewirings) increases, but also the number of coalition changes. Moreover, as we increase the probability of rewiring, the number of partner changes proportionally increases. However, coalition instability is lower than rewiring changes.

(3) Analysis of coalition agents' behavior.

- *Trading strategies.* As resource availability increases, coalition agents behaviors both against outsiders and insiders (coalition mates) varies. We have observed three scenarios: *low cooperation* (less than 20% of resources), where coalition

agents do not trade with outsiders, but mostly do it with insiders (around 75% of the agents behave as traders with coalition-mates); *no dominant strategy* (percentage of resources is around 20%), where half of coalitions behave as possessors with outsiders, while the other half behave as traders; and *high cooperation* (beyond 20% of resources), where trading becomes the dominant strategy against outsiders.

- *Rewiring behavior.* As the rewiring probability increases, the number of links that agents establish with coalition-mates increases. Moreover, agents behaving as traders with outsiders increase their number of links.

## 5. CONCLUSIONS

In this paper we have presented a cooperation mechanism to help self-interested agents to establish sustained and successful cooperation in resource exchange environments. The mechanism is based on: (1) a game-based interaction model that includes the trading of resources; (2) a dynamic coalition formation mechanism that allows agents to decide whether to join or leave coalitions, and to collectively self-determine the inner and outer behaviors of a coalition (without the intervention of a leader); and (3) a partner switching (rewiring) strategy based on experiences acquired in previous interactions to help agents to avoid defective behaviors.

We empirically analyzed our mechanism to quantify its benefits with respect to alternative cooperation mechanisms akin to others appearing in the literature. We observed that our mechanism allows agents to obtain higher payoffs, ranging from 15% to 30%, than other cooperation mechanisms. Furthermore, we found that these benefits depend on the availability of resources in the environment. Thus, the larger the number of resources, the larger the payoffs that agents are expected to obtain.

The benefits of our cooperation mechanism stem from the fact that rewiring has a positive effect when coupled with coalition formation. Thus, we empirically showed that rewiring helps agents to obtain higher payoffs (compared to exclusively using coalition formation, the so-called *coalition-only* mechanism), independently of the network topology. Indeed, even a small rewiring frequency can lead to up to a 20% increase of payoffs with respect to coalition formation. In fact, we observed that the higher the rewiring frequency, the less and bigger the number of coalitions, namely the lower the clustering of the agent population. In fact, our mechanism leads to less than a half of the coalitions created by means of coalition-only. Therefore, the use of rewiring has a compacting effect on the coalition formation process: less and bigger coalitions. Since agents within the same coalition cooperate, having less and bigger coalitions is bound to yield significant payoff benefits. The reason for this is that since agents are allowed to change their neighbors, they can choose to wire to agents that provide more benefits and to join bigger coalitions to be more efficient against non-cooperative behaviors. Together, the decrease in the number of coalitions, and the increase in their size, results in higher payoffs for agents.

Finally, we also analyzed agents' strategic behavior in the realm of our cooperation mechanism. Our main observation is that agents adapt their behavior depending on the availability of resources. In fact, in low cooperation scenarios, i.e., when there is less than 20% of resources, the dominant strategy with outsiders is to defect. When the percentage of resources is around 20%, there is a transition in the behavior of coalitions with outsiders, and there is no dominant strategy with outsiders. Beyond 20% of resources, coalition agents progressively become more and more cooperative with outsiders. Thus a medium or large availability of resources allows agents to perform more trades, obtaining higher benefits by cooperating in this way. To summarize, trading, i.e., cooperating, becomes the dominant strategy against outsiders.

Regarding future work, we plan to improve our model by developing new mechanisms to cope with agents with malicious behavior and/or free-riders, since we have seen that our approach is not able to successfully discriminate them. Moreover, we will also focus on studying sustainability of resources to analyze how it may influence the benefits and convergence of the approach. As we have seen that self-organization approaches and techniques used in [Pitt et al. 2012a] seem to be applicable to coalition formation, we plan to study the similarities between their approach and ours, to see how our results can contribute their problem, and vice-versa. Finally, we also plan to study different fair division techniques to see its influence in our approach.

## ACKNOWLEDGMENTS

The first author thanks the grant Formación de Profesorado Universitario (FPU), reference AP2010-1742. J.Ll.A. and J.A.R-A are partially funded by projects EVE (TIN2009-14702-C02-01), AT (CSD2007-0022), COR (TIN2012-38876-C02-01), MECER (201250E053), and the Generalitat of Catalunya grant 2009-SGR-1434

## REFERENCES

- ADAMIC, L.A. AND HUBERMAN, B.A. 2000. Power-Law distribution of the world wide web. *Science* 287, 5461, 2115.
- AUMANN, R. J. 1959. Acceptable points in general cooperative n-person games. In *Contribution to the theory of game IV, Annals of Mathematical Study* 40, R. D. Luce and A. W. Tucker, Eds. University Press, 287–324.
- AXELROD, R. 1997. *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration* 1st Printing Ed. Princeton University Press.
- AXELROD, R. M. 1984. *The evolution of cooperation*. Basic Books, New York.
- AZIZ, H., BRANDT, F., AND SEEDIG, H. G. 2011. Stable partitions in additively separable hedonic games. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1. AAMAS '11*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 183–190.
- BACHRACH, Y. AND ROSENSCHEIN, J. S. 2008. Coalitional skill games. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems - Volume 2. AAMAS '08*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1023–1030.
- BATAGELJ, V. AND MRVAR, A. 2003. Pajek - analysis and visualization of large networks. In *Graph Drawing Software*. Vol. 2265. 77–103.
- BAZZAN, A., PELETEIRO, A., AND BURGUILLO, J. 2011. Learning to cooperate in the Iterated Prisoner's Dilemma by means of social attachments. *J. Braz. Comp. Soc.* 17, 3, 163–174.
- BRANDT, F., CONITZER, V., AND ENDRISS, U. 2013. Computational social choice. In *Multiagent Systems*, G. Weiss, Ed. MIT Press, 213–283.
- BURGUILLO, J. AND PELETEIRO, A. 2010. Ownership and trade in spatial evolutionary memetic games. In *Proceedings of the 11th international conference on Parallel problem solving from nature: Part I. PPSN'10*. Springer-Verlag, Berlin, Heidelberg, 455–464.
- BURGUILLO-RIAL, J. 2009. A memetic framework for describing and simulating spatial prisoner's dilemma with coalition formation. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1. AAMAS '09*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 441–448.
- CHALKIADAKIS, G., ELKIND, E., MARKAKIS, E., POLUKAROV, M., AND JENNINGS, N. 2010. Cooperative games with overlapping coalitions. *Journal of Artificial Intelligence Research (JAIR)* 39, 179–216.
- CHEVALEYRE, Y., ENDRISS, U., LANG, J., AND MAUDET, N. 2007. A short introduction to computational social choice. In *SOFSEM (I) (2007-09-13)*, J. van Leeuwen, G. F. Italiano, W. van der Hoek, C. Meinel, H. Sack, and F. Plasil, Eds. Lecture Notes in Computer Science Series, vol. 4362. Springer, 51–69.
- DORAN, J. E., FRANKLIN, S., JENNINGS, N. R., AND NORMAN, T. J. 1997. On cooperation in multi-agent systems. *The Knowledge Engineering Review* 12, 309–314.
- EGUILUZ, V. M., ZIMMERMANN, M. G., CELA-CONDE, C. J., AND SAN MIGUEL, M. 2005. Cooperation and the emergence of role differentiation in the dynamics of social networks. *Am. J. Sociol.* 110, 4, 977.
- FU, F., TARNITA, C. E., CHRISTAKIS, N. A., WANG, L., RAND, D. G., AND NOWAK, M. A. 2012. Evolution of in-group favoritism.

- FU, F., WU, T., AND WANG, L. 2009. Partner switching stabilizes cooperation in coevolutionary prisoner's dilemma. *Phys. Rev. E* 79, 3, 036101.
- GRIFFITHS, N. AND LUCK, M. 2010. Changing neighbours: improving tag-based cooperation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1 - Volume 1*. AAMAS '10. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 249–256.
- GROSS, T. AND BLASIUS, B. 2008. Adaptive coevolutionary networks: a review. *J. R. Soc. Interface* 5, 20, 259–271.
- HOGG, T. 1995. Social dilemmas in computational ecosystems. In *Proc. of the 14th Intl. Joint Conf. on Artificial Intelligence (IJCAI95)*. Morgan Kaufmann, San Mateo, CA, 711–716.
- JACKSON, M. O., DEMANGE, G., GOYAL, S., AND NOUWEL, A. V. D. 2003. A survey of models of network formation: stability and efficiency. In *In Group Formation in Economics: Networks, Clubs and Coalitions*. Cambridge University Press.
- KNIESBURGES, S., KOUTSOPOULOS, A., AND SCHEIDELER, C. 2012. A self-stabilization process for small-world networks. In *Proceedings of the 26th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. to appear.
- LANGER, P., NOWAK, M., AND HAUERT, C. 2008. Spatial invasion of cooperation. *Journal of Theoretical Biology* 250, 4, 634–641.
- MAYNARD SMITH, J. AND PRICE, G. R. 1973. The logic of animal conflict. *Nature* 246, 5427, 15–18.
- NARENDRA, K.S. AND THATHACHAR, M. A. L. . 1989. *Learning automata: an introduction*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.
- NEWMAN, M. E. J. 2003. The structure and function of complex networks. *SIAM Review* 45, 2, 167–256.
- NGUYEN, D.-T. AND ISHIDA, Y. 2009. Spatial dilemma strategies of intelligent agents: Coalition formation in environmental game. In *Proceedings of the 2009 International Conference on Knowledge and Systems Engineering*. KSE '09. IEEE Computer Society, Washington, DC, USA, 126–129.
- NOWAK, M. AND MAY, R. 1992. Evolutionary games and spatial chaos. *Nature* 359, 6398, 826–829.
- NOWAK, M. A. AND MAY, R. M. 1993. The spatial dilemmas of evolution. *International Journal of Bifurcation and Chaos (IJBC)* 3, 1, 35–78.
- OSTROM, E. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Political Economy of Institutions and Decisions. Cambridge University Press.
- PACHECO, J. M., TRAUlsen, A., AND NOWAK, M. A. 2006. Coevolution of strategy and structure in complex networks with dynamical linking. *Physical Review Letters* 97, 25, 258103+.
- PASTOR-SATORRAS, R. AND VESPIGNANI, A. 2001. Epidemic dynamics and endemic states in complex networks. *PHYS.REV.E* 63, 066117.
- PELETEIRO, A., BURGUILLO, J., AND BAZZAN, A. 2011. Emerging Cooperation in the Spatial IPD with Reinforcement Learning and Coalitions. In *Intelligent Decision Systems in Large-Scale Distributed Environments. Studies in Computational Intelligence Series*. Springer London, Limited.
- PELETEIRO, A., BURGUILLO, J., AND BAZZAN, A. 2012. How Coalitions Enhance Cooperation in the IPD over Complex Networks. In *3rd Brazilian Workshop on Social Simulation*. IEEE Proceedings.
- PERC, M. AND SZOLNOKI, A. 2010. Coevolutionary games—a mini review. *Biosystems* 99, 2, 109–125.
- PITT, J., SCHAUMEIER, J., AND ARTIKIS, A. 2011. The axiomatisation of socio-economic principles for self-organising systems. *2012 IEEE Sixth International Conference on Self-Adaptive and Self-Organizing Systems 0*, 138–147.
- PITT, J., SCHAUMEIER, J., AND ARTIKIS, A. 2012a. Axiomatization of socio-economic principles for self-organizing institutions: Concepts, experiments and challenges. *ACM Trans. Auton. Adapt. Syst.* 7, 4, 39:1–39:39.
- PITT, J., SCHAUMEIER, J., BUSQUETS, D., AND MACBETH, S. 2012b. Self-organising common-pool resource allocation and canons of distributive justice. In *Self-Adaptive and Self-Organizing Systems (SASO), 2012 IEEE Sixth International Conference on*. 119–128.
- PUJOL, J. M., DELGADO, J., SANGESA, R., AND FLACHE, A. 2005. The role of clustering on the emergence of efficient social conventions. In *In IJCAI*. 965–970.
- RAHWAN, T. AND JENNINGS, N. 2008a. Coalition structure generation: Dynamic programming meets anytime optimisation. In *Proc 23rd Conference on AI (AAAI)*. 156–161.
- RAHWAN, T. AND JENNINGS, N. R. 2008b. An improved dynamic programming algorithm for coalition structure generation. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems - Volume 3*. AAMAS '08. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1417–1420.

- RAHWAN, T., MICHALAK, T., AND JENNINGS, N. R. 2012. A hybrid algorithm for coalition structure generation. In *26th Conference on Artificial Intelligence (AAAI-12)*. 1443–1449.
- RAHWAN, T., RAMCHURN, S., JENNINGS, N., AND GIOVANNUCCI, A. 2009. An anytime algorithm for optimal coalition structure generation. *Journal of Artificial Intelligence Research (JAIR)* 34, 521–567.
- REKA, A. AND BARABÁSI, A. 2002. Statistical mechanics of complex networks. *Rev. Mod. Phys.* 74, 47–97.
- SALAZAR, N., RODRIGUEZ-AGUILAR, J. A., ARCOS, J. L., PELETEIRO, A., AND BURGUILLO-RIAL, J. C. 2011. Emerging cooperation on complex networks. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*. AAMAS '11. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 669–676.
- SANDHOLM, T., LARSON, K., ANDERSSON, M., SHEHORY, O., AND TOHMÉ, F. 1999. Coalition structure generation with worst case guarantees. *Artif. Intell.* 111, 1-2, 209–238.
- SEO, Y.-G., CHO, S.-B., AND YAO, X. 1999. Emergence of cooperative coalition in nipa game with localization of interaction and learning. In *Evolutionary Computation, 1999. CEC 99. Proceedings of the 1999 Congress on*. Vol. 2. 3 vol. (xxxvii+2348).
- SEO, Y.-G., CHO, S.-B., AND YAO, X. 2000. Exploiting coalition in co-evolutionary learning. In *Evolutionary Computation, 2000. Proceedings of the 2000 Congress on*. Vol. 2. 1268–1275 vol.2.
- SERVICE, T. C. AND ADAMS, J. A. 2011. Constant factor approximation algorithms for coalition structure generation. *Autonomous Agents and Multi-Agent Systems* 23, 1, 1–17.
- SHEHORY, O. AND KRAUS, S. 1993. Coalition formation among autonomous agents: Strategies and complexity. In *MAAMAW*. 56–72.
- SHEHORY, O. AND KRAUS, S. 1995. Task allocation via coalition formation among autonomous agents. In *Proceedings of the 14th international joint conference on Artificial intelligence - Volume 1*. IJCAI'95. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 655–661.
- SHEHORY, O. AND KRAUS, S. 1996. Formation of overlapping coalitions for precedence-ordered task-execution among autonomous agents. In *ICMAS-96*. 330–337.
- SHEHORY, O. AND KRAUS, S. 1998. Methods for task allocation via agent coalition formation. *Artificial Intelligence* 101, 1, 165–200.
- SHEHORY, O., SYCARA, K., AND JHA, S. 1998. Multi-agent coordination through coalition formation. In *Intelligent Agents IV Agent Theories, Architectures, and Languages*, M. Singh, A. Rao, and M. Wooldridge, Eds. Lecture Notes in Computer Science Series, vol. 1365. Springer Berlin, Heidelberg, 143–154. 10.1007/BFb0026756.
- SHROT, T., AUMANN, Y., AND KRAUS, S. 2010. On agent types in coalition formation problems. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1 - Volume 1*. AAMAS '10. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 757–764.
- TANIMOTO, K. 2002. Coalition formation interacted with transitional state of environment. In *Systems, Man and Cybernetics, 2002 IEEE International Conference on*. Vol. 6. 6 pp. vol.6.
- VOICE, T., RAMCHURN, S. D., AND JENNINGS, N. R. 2012. On coalition formation with sparse synergies. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*. AAMAS '12. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 223–230.
- WATTS, D. J. AND STROGATZ, S. 1998. Collective dynamics of 'small-world' networks. *Nature* 393, 6684, 440–442.
- WOOLDRIDGE, M. AND DUNNE, P. E. 2006. On the computational complexity of coalitional resource games. *Artif. Intell.* 170, 10, 835–871.
- YEE, K. 2003. Ownership and trade from evolutionary games. *International Review of Law and Economics* 23, 2, 183–197.
- ZIMMERMANN, M. G., EGUILUZ, V. M., AND SAN MIGUEL, M. 2004. Coevolution of dynamical states and interactions in dynamic networks.