Coalition Formation with Boundedly Rational Agents*

D. Fiaschi - P.M. Pacini
Department of Economics
University of Pisa
Via Ridolfi, 10
56124 Pisa (Italy)

September 11, 2000

Abstract

Many economic activities take place in groups, teams, clubs (coalitions for short). However, because of informational problems, the allocation of agents in coalitional structures may face social dilemmas, since people may cooperate in groups of limited size but, in larger coalitions defection may become the dominant strategy. We analyze the process of coalition formation in which agents' expectations evolve through repeated interactions in a large population setting. The selection strongly depends on agents' initial beliefs; the efficient coalition structure is reached starting from a very limited set of initial beliefs and we show that the agents' computational ability and learning speed crucially affect outcomes. While overall efficiency is an increasing function of agents' computational ability, an increase in agents' learning speed can have an ambiguous effect.

*We thank C. Bianchi, C. Casarosa, E.M. Cleur, L. Fanti, N. Salvadori, P. Vagliasindi and all participants to seminars delivered at the University of Pisa and Venice for helpful comments; none of them is responsible for what is written here. This work has been accomplished with MURST national funds (40%).

1 Introduction

Cooperative behaviour often emerges at a group, rather than social level; in many instances we observe the formation of independent groups, teams, clubs, cooperatives (coalitions for short) each of them persecuting the same goal (in turn provision of commodities, maximization of profits, raising of public funds, standards of behavior etc.). This behaviour has been mostly analyzed within the theoretical apparatus of cooperative games; however this approach has recently undertaken a substantial revision in order to explain the emergence of coalitional structures (**CS**) as the outcome of a bargaining process in which agents cannot be committed to binding pre-play agreements; therefore the **CS** formation game has been recast within the framework of non-cooperative games (pioneering works are [12] and [15], see also [16])[1]. Adopting this non-cooperative framework we analyze the problem of how a population of individuals, persecuting the same goals, structures in coalitions when agents and groups are subject to two competing forces: on one hand an increasing returns to scale technology that incentivates aggregation, on the other the non-monitorability of actions in formed coalitions that incentivates free-riding behaviours, thus reducing the incentive to form and act in groups[2]. In fact, as it will be shown in Section 2, this situation is compatible with the presence of a multiplicity Nash equilibria, in the sense that there are many, qualitatively different, partitions of society in cooperation groups where no agent has an incentive to deviate from his course of action. Traditionally this indeterminacy has been reduced appealing to a sound refinement proposed by [7], i.e. *coalition proofness.* However coalition-proofness requires the possibility of a pre-play communication stage, in which agents can agree at no cost to correlate their strategies (for example a simultaneous exit from a coalition to form a new one), a stage that is difficult to justify when the population is large, as in the case we are going to study. To overcome this difficulty, we take another point of view and substitute the implicit pre-play communication phase with an explicit dynamic process in which anonymous agents repeatedly interact without the possibility of making jointly agreed

[1] This type of games well adapts to represent situations concerning the decentralized supply of public goods (e.g. see [17]), the formation of cooperative firms (see [12]) or multilateral bargaining; more generally they can be applied to all those situations in which there is a problem of coordination and imperfect information (as to this point see also [13]).

[2] For example, take the case of the formation of work cooperatives (research groups as well). The larger the group the greater the potential output because of the integration of different competencies and the cutting of administrative costs; but the larger the group the more a single has an incentive to shirk, benefitting of collective output and saving on the individual effort directed to the coalitional goal.

deviations. At any stage any agent plays the strategy that is best for him, *given his beliefs* about the behaviour of the others. The outcomes of individual strategies provide the information on the basis of which beliefs are revised, thus moving to a further stage of the game. An equilibrium is found when a resting state is encountered in which beliefs and actions no further change. In this way we try to provide an answer to the following questions: when the population is so large that it is not reasonable to assume that pre-play agreements can be reached, does a system in which agents act rationally find an equilibrium position? And, provided that it does, what are the characteristics of the equilibrium positions? (see also [21])

Our main finding is that agents play the strategies sustaining a coalition proof Nash equilibrium (CPNE henceforth) only if population start out with a priori beliefs that are close to those sustaining coalition-proof outcomes; in other words, in this dynamic setting coalition-proof positions can be reached only if a large fraction of the population initially share the conviction that all other people will act and cooperate in large coalitions; the role of learning will be just to size down the dimension of the equilibrium coalitions to that consistent with individual incentives to cooperate. On the contrary, if initial beliefs are dispersed and there is not such "optimistic" initial state of expectations, the interaction among people will drive the system toward equilibrium positions that are still Nash situations, but society is divided into a larger number of groups than that corresponding to CPNE. To explanation this fact consider a population in which there are also "less optimistic" agents who believe that cooperation can be sustained only in small groups. Then the latter ones will free-ride in large coalitions: this behaviour will induce the revision of expectations for the "optimistic" agents who, in turn, will be no longer ready to form large coalitions and accept the free-riding behaviour of the "less optimistic" players, provided the latter ones cannot be discriminated. This will determine a downward revision of the expectations on the cardinality of the maximal coalition capable of sustaining full cooperation, so that, eventually, any equilibrium will be characterized by a greater number of smaller coalitions than those would be observed were the population made up only by "optimistic" agents. In particular we will show that there are two main sets of resting points (attractor sets in the following); the first one is made up of coalitional structures "close" to CPNE, i.e. there is a very limited number of agents refusing aggregation, while the large part of the population behave according to the CPNE prescriptions. The other attractor set is characterized by coalitional structures in which coalitions are, on the average, of smaller size, so that, given the available technology, the overall result is certainly less efficient. Moreover the latter shows a greater basin of attraction than the first one.

The analysis is performed by means of computational experiments reproducing an artificial society in which adaptive and heterogeneous[3] agents interact. Their behaviour results from two simultaneous learning processes[4]: the first one deals with the revision of beliefs about the possible realizations of an action, so that agents learn what will be the plausible outcome of a strategy conditional on others' behaviour; this revision of expectations will take the form of a simple reinforcement learning process. The second one deals with the choice of the course of action: given a certain set of beliefs, an agent finds his best strategy through a process of *learning* and *imitation* on the set of possible strategies; this decision process is modelled by means of *genetic algorithm* (GA)[5] which tends to capture the idea that agents, when called upon to make a choice in a complex environment, do not make explicit optimization, but rather operate on a set of rules that they continuously modify reacting to the effects of their own behaviour. In this framework we tested the robustness of our results to various parametrizations of both learning processes; we find that the overall allocation efficiency is an increasing function of agents' computational ability, while an increase in agents' learning speed can have a ambiguous effects (see Section 4.3).

For both the problem and the techniques used for the analysis, this work can be linked with those contributions analyzing the auto-organization in economies where agents take decisions on the basis of a limited knowledge of the environment: for example [11] studies the formation of coalitions in an economy in which any agent can deal only with a limited number of other agents and learning takes place by means of neural nets, while [24] analyzes the auto-organization in an economy in which agents must coordinate to elaborate at best the information they have and use GA to learn which is

[3]Here heterogeneity refers to agents' initial beliefs on others' actions. Indeed there could be a further source of heterogeneity (as indicated by [10]), i.e. the different sets of strategies with which agents start playing the game.

[4]The learning process is a crucial aspect in the modelling of any evolutionary game (for an excellent review see [23]). A distinguishing feature of the present approach is that the dynamics of the system is not driven by the payoffs that strategies receive when played, but by the configuration of beliefs that change over time as new realizations of the system occur; these affect the expected payoffs and, by this way, the strategies actually played.

[5]The use of GA to represent individual behaviour has been motivated by important contributions from the theory of cognitive processes (see [19]). They have been increasingly applied in economics: as to the applications to game theory the pioneering work is [6]. The extension to evolutionary games is recent: [9] provides some examples, [25] highlights how information affects the evolution of players' strategies in a prisoner's dilemma game. Most contributions using GA in economic analysis focus on macroeconomic models with rational expectations and multiple equilibria (see [10], [22], [2], [3], [4] and [10]), but some authors also tackle explicitly the decision problem (see [5], [18] and [20]).

the most efficient structure.

The paper is organized as follows: Section 2 describes the basic characteristics of the model, the strategies available to agents and the kind of beliefs they initially have about the state of the world. Section 3 describes the dynamic process of coalition formation and how, in this process, expectations are revised and decisions are taken by means of the GA procedure. Section 4 describes the calibration of the artificial experiments, the results and their main characteristics, providing also some insights on how the results change when we change the basic parameters affecting agents' learning process. Conclusions close the paper. Some technical material is relegated into the Appendices.

2 The model

The basic characteristics of the model are the following[6]: there is a population $\Im$ of I agents indexed by i. Agents are identical in all physical characteristics and are endowed with 1 unit of time that they can use either working ($l_i = 1$) or as leisure time ($l_i = 0$). They receive utility from the consumption of a commodity y and leisure time, according to the following utility function

$$U_i\left(y_i, l_i\right) = y_i + \left(1 - l_i\right) \cdot \omega,$$

where ω measure the pleasure of not working.

Agents can form coalitions, but an agent can participate into one and only one coalition. A coalition is a group of agents that agree to share the output they produce by means of the total labour input that they provide. Labour is the only productive input and the technology for the production of the coalitional output Y is given by the production function $Y = L^{\alpha}$, where L is the number of labour units and α is assumed to be greater than 1. The production within a coalition has no external effect on the production of other coalitions; commodity Y deteriorates in a single period.

We assume that the agreement within a group entails an *equal sharing* distribution of the coalitional output, i.e. in a coalition of N agents anyone receives an amount $y_i = \frac{Y}{N}$ of the produced output Y. While participation into a group is publicly observable, the contribution of the individual working time to the coalitional production process cannot be monitored, so that defection ($l_i = 0$) cannot be punished.

Within this setting agent has to make two decisions: (i) which coalition to participate into and (ii) which action to perform in a coalition, once formed.

[6]For the complete description of the model and proofs of statements see [13]

The first action determines the formation of a **CS** σ, i.e. a partition of the population $\Im$ in coalitions S_k, while the second determines the individual payoffs within a coalition S_k.

Strategies In this framework a strategy for an agent i must concern (i) the formation of a coalition and (ii) the action to perform in a coalition once formed. Therefore any strategy θ_i is made up of two components, i.e. $\theta_i = \{\theta_i', \theta_i''\}$: θ_i' is a *signal* indicating the maximum cardinality of the coalition that i is ready to form[7]; θ_i'' is a complete contingent plan indicating the action l_i that i is willing to take conditional on the cardinality of the coalition he may happen to belong to. As an example a strategy θ_i can take the form

$$\theta_i = \{\theta_i', \theta_i''\} = \{N, [1, 1, 1, 0, \ldots]\}$$

meaning that agent i is ready to participate into any coalition of cardinality at most N and cooperate ($l_i = 1$) in all coalitions with cardinality less or equal to 3, while he will not cooperate ($l_i = 0$) in coalitions of greater cardinality. Finally, we assume that agents cannot play mixed strategies.

Coalition formation Once strategies are announced, agents randomly match. If two agents i and j match and $\min\left\{\theta_i', \theta_j'\right\} \geq 2$ then the coalition $S = \{i, j\}$ forms; this coalition is ready to accept another (randomly chosen) agent h provided $\min\left\{\theta_i', \theta_j', \theta_h'\right\} \geq 3$ otherwise h will be the first member of a new coalition S' and so on. A **CS** σ is obtained when $\Im$ is partitioned in groups S in such a way that $\min\{\theta_i'\}_{i\in S} \geq |S|$ ($\forall S \in \sigma$) and there is no couple of groups S and S' in σ such that $\min\left\{\min\{\theta_i'\}_{i\in S}, \min\{\theta_i'\}_{i\in S'}\right\} \geq |S| + |S'|$, i.e. a **CS** σ is formed whenever no agent is compelled to participate into a coalition of greater cardinality than the one he is willing to accept and no two groups are compelled to remain separated when they could join without the objection of any participant.

Equilibria A configuration of strategies $\theta^* = \{\theta_i^*\}_{i\in\Im}$ is a Nash equilibrium of the game if it gives rise to a **CS** σ^* and a corresponding profile of actions l^* such that no agent has an incentive to deviate from his strategy. Indeed this game may have many Nash equilibria and the following example can give an intuition of this fact:

[7] This is motivated by the consideration that agents are identical in all physical respects and differ only as to unobservable characteristics, so that their identity does not matter for the formation of coalitions. The fact that θ_i' is the *maximal* acceptable cardinality implies that an agent, signalling θ_i', refuses to participate in all coalitions larger than θ_i' but is ready to belong to coalitions of smaller cardinality.

Example 1 *Suppose that $I = 4$, $\omega = 0.635$ and $\alpha = 1.428$. The individual payoffs from cooperation and defection are reported in the following Tables, where rows represent the cardinality of a possible coalition and columns the number of other cooperators in the corresponding coalition*

	Cooperation			
	0	*1*	*2*	*3*
1	1	.	.	.
2	0.5000	1.3454	.	.
3	0.3333	0.8969	1.6003	.
4	0.2500	0.6727	1.2002	1.8100

	Defection			
	0	*1*	*2*	*3*
1	0.6350	.	.	.
2	0.6350	1.1350	.	.
3	0.6350	0.9683	1.5319	.
4	0.6350	0.8850	1.3077	1.8382

Clearly the grand coalition with full cooperation can never form as a Nash equilibrium, because everyone has an incentive to defect. However each of the following strategy profiles is a Nash equilibrium

$$\begin{aligned} \forall i, \quad \theta_i^* &= \{1, [1,0,0,0]\} \\ \forall i, \quad \theta_i^{**} &= \{2, [1,1,0,0]\} \\ \forall i, \quad \theta_i^{***} &= \{3, [1,1,1,0]\} \end{aligned}$$

They give rise to the following partitions of society in groups characterized by full cooperation

$$\begin{aligned} \theta^* &\rightarrow \sigma^* = \{S_k\}_{k=1}^{4}, \quad |S_k| = 1 \\ \theta^{**} &\rightarrow \sigma^{**} = \{S_k\}_{k=1}^{2}, \quad |S_k| = 2 \\ \theta^{***} &\rightarrow \sigma^{***} = \{S_1, S_2\}, \quad |S_1| = 3, \quad |S_2| = 1 \end{aligned}$$

Indeed any partition of society in groups $\{S_1, \ldots, S_K\}$ such that $1 \leq |S_k| \leq \bar{N}(\alpha)$ where $\bar{N}(\alpha)$ is the cardinality of the maximal coalition in which cooperation is incentive compatible for all its members[8], is a Nash equilibrium for a suitable configuration of strategies; for example the **CS** $\sigma' = \{S'_1, \ldots, S'_K\}$, $1 \leq |S_k| \leq \bar{N}(\alpha)$, is certainly a Nash equilibrium if i's strategy is $\theta_i = \{\theta'_i, \theta''_i\} = \{|S'_k|, [1, \ldots, 1, 0, \ldots]\}$, $S_k \ni i$, where the last 1 in the conditional action part of θ_i is in the $|S'_k|^{th}$ position.

[8]Provided $\frac{1}{2} < \omega < 1$, $\bar{N}(\alpha)$ is the integer part of the value of N solving the equation $N^\alpha - (N-1)^\alpha - \omega \cdot N = 0$. $\bar{N}(\alpha)$ is monotonically increasing in α, so that all agents merge and cooperate in the grand coalition provided $\alpha \geq \alpha^* > 1$, where α^* solves $I^{\alpha^*} - (I-1)^{\alpha^*} - \omega \cdot I = 0$. Furthermore, since $\bar{N}(\alpha) < 2$ for $\alpha \leq 1$, it follows that increasing returns are a necessary condition in order to observe cooperation in groups of at least 2 agents. Finally, for $\alpha \in (1, \alpha^*)$ cooperation can emerge just in coalitions that are proper subsets of $\Im$; this is the case we will deal with in the rest of the paper.

Among all Nash equilibria the **CS** in which all, but possibly one, coalitions have cardinality $\bar{N}(\alpha)$, is particularly interesting because it implies the maximum aggregate output and furthermore it is a *coalition-proof Nash equilibrium* (see [13]). However the implementation of this refinement presumes a lot of communication and coordination capacities on the side of agents. To see this take the Nash equilibrium situation in which every agent plays the strategy $\theta_i^* = \{1, [1, 0, 0, 0]\}$ and the CPNE profile in which, for every i, $\theta_i^{***} = \{3, [1, 1, 1, 0]\}$; to reach the second from the first at least three agents has to correlate their strategies and mutate both the maximal acceptable cardinality and the conditional action part of their strategies. While this seems plausible in small groups, where communication can take place easily and at no cost, it seems less justifiable in large societies where the search for precommitments may be very resource expensive. Therefore the problem remains open about what are the most plausible outcomes of the social interaction out of the multiplicity of Nash equilibria, when the pre-play communication phase is not allowed.

In this paper we analyze this equilibrium selection problem by adopting an evolutionary approach. We suppose that agents play repeatedly the game; at each stage they choose their strategies on the basis of their beliefs on the other agents' behaviour, having as time-horizon only one period (i.e. they play a series of one-shot games). Then, at the end of every period, agents revise their beliefs on the basis of experience. While the fact that agents will finally play Nash equilibria is an expected outcome, our focus will be on the question on which Nash equilibrium agents will play. In this we will pay attention to key factors, as initial beliefs and shape of learning processes, driving the selection among the possible (Nash) equilibria.

The analysis is performed by computational experiments; the next Section describes the details of the artificial setting we use for the simulations.

3 The design of the simulation

The basic idea behind the simulation is to represent the evolution of the game as a sequence of periods. At the beginning of each period every agent plays a strategy stating the maximum cardinality of the coalition he intends to belong to and the action he will take, conditional upon the effective cardinality of the realized coalition. The choice of a strategy is the outcome of a maximizing process *conditional on the current period beliefs* about the actions of the others; the decision process is implemented by *GA*, that, as we will see, provide a very intuitive way to model learning and emulation, two crucial aspects of any decision process.

The strategies $\theta = (\theta_1, ..., \theta_I)$ played by the agents and the random matching process described in the previous Section determine the **CS** σ and the agents' actions $l = (l_1, ..., l_I)$ of the current period. Given σ and l, every player uses this new information to upgrade his beliefs and receives his payoff. This concludes the period. The simulation goes on until a persistent pattern in the **CS** emerges (see Appendix B for more details). In the following we describe the various components of this procedure in more details[9].

3.1 Initial beliefs and agents' type

In this game agents play their "best" strategy *given* their expectations on others' behaviour; according to the two parts of a strategy, agents have two forms of expectations:

1. the probability that a coalition of size N will form if one communicates his willingness to participate in coalitions of size at most S. This probability is indicated by $P_i(N|S)$ and we assume that, at the beginning, everyone believes that all coalitions of smaller size than S are equiprobable, i.e. $\bar{P}_i(N|S) = \frac{1}{S}$, $1 \leq N \leq S$, and $\bar{P}_i(N|S) = 0$, $N > S$, for all $S \in [1, I]$. This hypothesis corresponds to the case that agents do not know the rule governing the formation of coalitions, except for the fact that they are to be voluntary[10].

2. the probability that other people in a coalition will cooperate; in this respect, and to keep things as simple as possible, we distinguish only three different types of initial beliefs

    - *optimistic* agents: agents of this type assign high confidence to the fact that other people are cooperators; more precisely they believe that any other agent is a cooperator with probability $p = 0.9$ irrespective of the size of the coalition he may be into. Therefore the initial probability $\bar{Q}_i(n|N)$ of finding n other cooperators in a coalition of cardinality N is the value at n of the binomial distribution with probability 0.9 and $N - 1$ trials.
    - *mildly optimistic* agents: they start playing with a lower confidence about the cooperative attitude of the other agents; for them the probability that any other agent will be a cooperator,

[9] The software used for the simulation is available from the authors upon request.

[10] We performed our analysis also for the case in which agents have heterogeneous beliefs on the coalition cardinality and the results do not change. This shows that the cooperation beliefs are the key factor in deciding which strategy has the highest expected payoff.

irrespective of the size of the coalition, is $p = 0.6$, i.e. the initial probability $\bar{Q}_i(n|N)$ of finding n other cooperators in a coalition of cardinality N is the value at n of the binomial distribution with probability 0.6 and $N-1$ trials.

- *pessimistic* agents: agents of this type assign a very low probability to the fact that anyone else is a cooperator, i.e. for them we set $p = 0.3$.

It is trivial to observe that the more optimist an agent is (i.e. the higher the value of p) the greater is the cardinality of the maximal coalition in which the corresponding agent is ready to cooperate (naturally always within the limit of $\bar{N}(\alpha)$)[11]. As an example take the situation described in the example 1; the following table reports the expected payoffs for eight different strategies related to the three different types of initial beliefs:

		Initial Beliefs		
		Opt.	M. Opt.	Pess.
Strategies	$\{1, [1,0,0,0]\}$	1	1	1*
	$\{1, [0,0,0,0]\}$	0.635	0.635	0.635
	$\{2, [1,1,0,0]\}$	1.1304	1.0036*	0.8768
	$\{2, [1,0,0,0]\}$	1.0425	0.9675	0.8925
	$\{3, [1,1,1,0]\}$	1.2406*	1.0224	0.8126
	$\{3, [1,1,0,0]\}$	1.2274	1.0417	0.8698
	$\{4, [1,1,1,1]\}$	1.3379	1.0466	0.7739
	$\{4, [1,1,1,0]\}$	1.3505	1.0810	0.8356

We can see that, for the optimists, the maximal coalition that can form with full cooperation is of size 3; indeed, even if to play a strategy with a maximum acceptable coalition of 4 has a greater expected payoff, however the

[11] The values of the probabilities determining the degree of optimism are appropriately chosen in order to make the optimist to sustain cooperation in the 3 person coalition, the mildly optimist in the 2 person coalition and the pessimist to sustain cooperation only if he is alone. The following table reports in the second row the minimum values of the probability with which an agent in a coalition cardinality from 1 to 4 (first row) should expect that everyone else in his coalition will cooperate, in order to cooperate himself:

$\|S\|$	1	2	3	4
p	0	0.39	0.77	1

.

This means that cooperation is always the best action for an agent in the singleton coalition; in a coalition of 2 persons, an agent cooperates only if the probability with which he expects that the other cooperates is at least 0.39 and so on. Defection is always the best action when an agents finds himself in a coalition of cardinality greater than 3.

dominant strategy in that situation would be to defect. In the same manner the mildly optimists will sustain full cooperation in coalition of cardinality 2; finally pessimists will cooperate only in the singleton coalition. In other words, optimistic beliefs can sustain a *strongly associative* (SA) behaviour, mildly optimistic beliefs can sustain a *weakly associative* (WA) behaviour and pessimistic beliefs sustain only a *non associative* (NA) behaviour. Starting from this consideration, in the sequel we will term SA those that are ready to aggregate and sustain cooperation in coalitions of at least 3 persons, WA those that are ready aggregate and sustain cooperation in coalitions of 2 persons and NA those that are ready to form just singleton coalitions. One of the purposes of the simulation is just to map any configuration of initial beliefs to a configuration of final behaviours, in order to see how interaction and learning modify beliefs and consequently final outcomes.

3.2 Revision of beliefs

Individual expectations are revised through time. Starting from prior beliefs, posterior distributions are formed taking into account observations; observations are relative to the local experience of an agent, i.e. he can know only the cardinality of the coalition he happens to be into and the profile of actions of the other members. We assume that the mechanism governing the process of revision of expectations is of a very simple type and takes the characteristics of a *reinforcement learning*. In this we assume that agents are *naive* in the sense that

1. the realization of a coalition (and its cardinality) is regarded as contingent upon the strategy an agent plays, so that the realization of a coalition of small cardinality when the strategy signalled the willingness to accept only coalitions of small size does not affect the probability that an agent assigns to the same coalitional size as a consequence of communicating the willingness to accept larger coalitions *and*

2. the realization of a certain level of cooperation is contingent upon the cardinality of the realized coalition, so that the realization of a low level of cooperation in a large group does not affect the probability assigned to the same level of cooperation in coalitions of different size.

To formalize this process of beliefs revision, consider the expectations about the cardinality of a coalition[12] and let $H_i^t(N|S)$ be the relative frequency with which agent i observed the formation of a coalition of cardinality

[12] The same analysis could be performed as to the beliefs about the expected number of cooperators.

N up to period t as a consequence of playing a strategy θ in which $\theta_i' = S$; then for any $t > 0$

$$\begin{aligned} P_i^{t+1}(N|S) &= \begin{cases} (1-\delta) \cdot P_i^t(N|S) + \delta \cdot H_i^t(N|S) & \text{if } \theta_i'^t = S \\ P_i^t(N|S) & \text{otherwise} \end{cases} \\ P_i^0(N|S) &= \bar{P}_i(N|S) \text{ given,} \end{aligned} \tag{1}$$

where δ is a parameter measuring the importance of experience in the formation of beliefs.

Equation (1) states that agents revise the prior on the possible events ensuing from S on the basis of the relative frequencies with which they occurred in the past, including last period observation, but the same last period observation is not used to modify the beliefs concerning the possible outcomes of different strategies.

To better understand the properties of this learning mechanism, assume that an agent i plays always the same strategy, so that he is always ready to accept coalitions of the same cardinality S; then we have that

$$\begin{aligned} P_i^{t+1}(N|S) &= (1-\delta) \cdot P_i^t(N|S) + \delta \cdot \frac{T(N)}{t} \\ P_i^0(N|S) &= \bar{P}_i(N|S) \text{ given} \end{aligned}$$

where $T(N)$ is the number of times a coalition of cardinality N occurred up to period t. By this we get that

$$P_i^t(N|S) = (1-\delta)^t \cdot \bar{P}_i(N|S) + \delta \cdot \sum_{s=1}^{t} (1-\delta)^{t-s} \cdot \frac{T(N)}{S}$$

i.e. there is a progressive decay of a priori probability $\bar{P}_i(N|S)$ at a rate $-\delta$ and, in the limit, the individual beliefs tends to the observed frequencies. We stress that, even in the case in which all agents would share the same initial beliefs, since every agent has different information about the history, the individual beliefs, at least in the first period, may become heterogeneous.

3.3 The decision process

GA is the engine by which strategies are selected and reproduced to arrive to a strategy that is "optimal" given individual beliefs; in this process also new strategies are created and evaluated. This algorithm seems well suited to our purposes because it offers a very intuitive way to model the decision process typical in the evolutionary approach, where trial-and-error learning and imitation are two crucial aspects[13].

Building on the idea of natural selection, GA start working on a set of candidate strategies for a given period and select the strategies with the highest fitness, calculated as the expected payoff given the beliefs on others' strategies, and "*recombine*" their single components (*building blocks*) to produce new ones. An intuitive interpretation of these building blocks is to consider a strategy as the result of different components; for example a component could be the action for a coalition of a certain cardinality; if this action is particularly efficient (and this is measured by the fitness of string), then GA will tend to use this component in the formation of new strategies. The performance of GA is further aided by what has been termed *imitation*: in this context imitation means that an agent uses the observable part of the strategies played by others in his coalition as new building blocks (genetic material) to produce better strategies. In the following we briefly describe the working of the GA procedure.

Encoding The first step is to represent the strategy in a way that can be handled by GA. This is done by coding a strategy in a binary alphabet as a string of bits; the first indicates the maximum cardinality of the coalition agent intends to belong to (this encodes θ_i') and the second the contingent plan of actions (θ_i'') conditional on the cardinality of possible coalitions. For example take an economy with 4 agents; in such a case a possible strategy is a string like the following:

$$\underbrace{\boxed{0}\boxed{1}}_{\text{maximum cardinality}} \quad \underbrace{\boxed{1}\boxed{1}\boxed{1}\boxed{1}}_{\text{conditional actions}}$$

By this strategy an agent intends to form a coalition with a maximum acceptable cardinality of 3 (0 1 in binary alphabet) and cooperate in every coalition he can belong to (the first three bits of the second substring are 1, which means cooperation)[14].

Therefore the set of rules (strategies) is a $J \times L$ binary matrix, where J is the number of strategies coexistent in the population (the "mind" of an agent) and L is the number of bits necessary to express a **CS** in binary form.

[13] Moreover, by a computational point of view, GA has been proved to be very efficient in searching for a solution in highly dimentional spaces, which is particularly useful in our case, where agents are looking for the best strategies in a solution space that, according to our encoding procedure (see later), is the $\left(I + \frac{\log(I)}{\log(2)}\right)$-dimentional hypercube.

[14] We use a string with a fixed length to simplify algorithm implementation, although, in this case, the 4^{th} bit in the second substring is not necessary (indeed the maximum cardinality for this strategy is equal to 3).

As it is well known, GA work sequentially and their procedure is made up by three basic steps: (i) selection, (ii) recombination and (iii) mutation. Each of them will be dealt with separately.

Selection In the biological evolution, the greater is the ability of a species to adapt and cope with the environment, the higher is the probability for it to survive and reproduce; similarly GA privilege those strings (strategies) with the highest fitness, giving them the highest chance to survive and reproduce. In order to model this mechanism we assume that J strings are drawn from the available population, where the probability of a string to be drawn is positively correlated to its fitness indexes $F_i^t(j)$[15]. In this way the highest is the fitness index the highest is the probability that the corresponding strategy is selected and passed to the next step of the procedure; conversely strategies with a low fitness index are candidate to be eliminated soon, since they proved inefficient.

Cross-over The mere selection of the fittest strategies serves the purpose to refine the set of current schemas but does not allow for the discovery of better ones. This further step is accomplished by recombining the building blocks of the selected proposals, as in the natural process of procreation and consequent exchange of genes. By "mixing" the building blocks (genes) of the fittest strategies we get new strings with an hopefully enhanced capacity of adaptation to the environment; such a mixing is called *cross-over*. There are several ways to model the crossing over of strings. We adopt the most commonly used in the literature: first a couple of strings is picked up at random from the set of selected strings (they are candidate to be parents) and, with a certain probability, they mix their genes. This means that each of them is partitioned into two substrings of length v and $L-v$ respectively, where v is a random integer drawn from an uniform distribution over the interval $[2, L-1]$; finally two substrings of equal length are interchanged and, by recombination, we get two new strings. For example take the pair of strings

1	1	0	1	1	0

1	0	1	0	1	0

and suppose that the random integer $v = 2$ is drawn; then each string is cut respectively into two substrings of length 2 and 4 and the two substrings of

[15] In our computational experiment, the Geometric Selection algorithm has been used; see Appendix A for details.

length 4 are interchanged

$$
\begin{array}{|c|c|}\hline 1 & 1 \\ \hline \end{array}\quad
\begin{array}{|c|c|c|c|}\hline 0 & 1 & 1 & 0 \\ \hline \end{array}
$$

$$\downarrow\uparrow$$

$$
\begin{array}{|c|c|}\hline 1 & 0 \\ \hline \end{array}\quad
\begin{array}{|c|c|c|c|}\hline 1 & 0 & 1 & 0 \\ \hline \end{array}
$$

Recombination of substrings finally yields the following two new strings (strategies)

$$
\begin{array}{|c|c|c|c|c|c|}\hline 1 & 1 & 1 & 0 & 1 & 0 \\ \hline \end{array}
$$

$$
\begin{array}{|c|c|c|c|c|c|}\hline 1 & 0 & 0 & 1 & 1 & 0 \\ \hline \end{array}
$$

The newly generated proposals are substantially different from their parents, although they are "based" on them.

Mutation A crucial element in any evolutionary process is chance. The presence of chance in our context is taken into account by adding a further step in the GA procedure in which every single bit in the set of strings is subjected (with a low probability) to a random mutation of its state. By this trick we avoid the lock-in phenomenon and we help search to escape from local inefficient optima.

Imitation A crucial aspect of any learning process is to emulate the most successful strategies played by the opponents. To make more realistic the analysis we consider that agent i can imitate only what he can observe, i.e. only the cardinality of his coalition, that provides new information on the maximum acceptable cardinality, and the actions *actually* played by the agents belonging to his coalition (not their full strategies, that are unobservable). Therefore he will replace some of the best candidate strategies with new ones obtained modifying the formers both in the maximal cardinality part and conditional action part so as to incorporate the new accruing information. For example if agent i is in a coalition of 3 agents and all others cooperate, he modifies some of his strategies so as to set to 3 the maximum acceptable cardinality and to 1 the third bit of the second substring. Therefore, if one of his strings was

$$
\begin{array}{|c|c|}\hline 1 & 1 \\ \hline \end{array}\quad
\begin{array}{|c|c|c|c|}\hline 1 & 1 & 0 & 0 \\ \hline \end{array},
$$

he could modify it according to observed realization as to become

$$
\begin{array}{|c|c|}\hline 0 & 1 \\ \hline \end{array}\quad
\begin{array}{|c|c|c|c|}\hline 1 & 1 & 1 & 0 \\ \hline \end{array}.
$$

Fitness The fitness of a string is nothing but its expected payoff conditional on the individual beliefs $P_i^t(.,.)$ and $Q_i^t(.,.)$ in the period in which it is to be evaluated. In particular, if an agent consider playing a strategy $\theta = \{\theta', \theta''\} = \{\theta', [\theta''_1, \ldots, \theta''_N, \ldots \theta''_I]\}$, $Q_i^t(n, N) \cdot P_i^t(N, \theta')$ is the probability with which i expects to find himself in a coalition of N people with other n and $Q_i^t(n, N) \cdot P_i^t(N, \theta') \cdot \left(\frac{(n+\theta''_N)^\alpha}{N} + (1 - \theta''_N) \cdot \omega \right)$ is the expected payoff in that situation. Therefore the expected payoff of playing θ in period t, i.e. the fitness $F_i^t(\theta)$ to the strategy $\theta = \{\theta', \theta''\}$ is given by

$$F_i^t(\theta) = \sum_{N=1}^{\theta'} \left[P_i^t(N|\theta') \cdot \sum_{n=0}^{N} Q_i^t(n|N) \cdot \left(\frac{(n + \theta''_N)^\alpha}{N} + (1 - \theta'_N) \cdot \omega \right) \right]$$

where θ''_N is the N^{th} component of the conditional action part of the strategy θ.

4 Computational experiments and results

Our computational experiments can be divided in two main stages; in the first we analyze how, setting the parameters regarding learning speed and computational ability to appropriate values, a population composed by agents with heterogeneous initial beliefs evolves through repeated interactions by learning and experimentation (see Section 4.2). In the second we analyze how our findings are robust to changes in the degree of computational ability and in the learning speed (see Section 4.3).

In any stage we run several simulations changing, for a given distribution of initial beliefs, the seed of random number. This is necessary to eliminate possible random disturbances, deriving from the fact that agents' initial set of candidate strategies are randomly generated and GA make use of random numbers. From another point of view, we can say that a further source of heterogeneity among agents is introduced by the differences in the initial set of candidate strategies (see [26]), so that it is necessary to run many simulations with different random numbers in order to eliminate this fact. We identify equilibria with positions of the system that do not show "significative" changes over appropriate number of periods. The presence of such *resting states* is captured by a stopping condition, that is detailed in Appendix B.

The next Section reports the value of parameters common to all simulations.

4.1 Parameterization

For simplicity we limit our attention to an economy with just 16 agents; even if the economy is so simple, the set of possible strategies is very large being made up by 2^{20} elements. As regarding to the parameters of model we set $\alpha = 1.428$ and $\omega = 0.635$; these values respect all the constraints of Section 2, i.e. $1 < \alpha < \alpha^*$ (in the present case the condition $I^\alpha - (I-1)^\alpha - \omega I = 0$ gives $\alpha^* = 1.6608$) and $\frac{1}{2} < \omega < 1$. Given these parameters, the cardinality of the greatest coalition capable of sustaining cooperation is $\bar{N}(\alpha) = 3$, so that every partition of 16 agents in coalitions with 3 or less agents with full cooperation are Nash equilibria. The CPNE is given by 5 coalitions of 3 agents and 1 of one agent; it is easy to check that this configuration implies maximum aggregate output, subject to the individual incentive constraints. Regards to GA, we set to 30 the number of strings forming the individual population of rules (that is $J = 30$), the parameter for the geometric selection to 0.8 (that is $p^{gs} = 0.8$, see Appendix A), the crossover probability to 0.6 and the mutation probability to 0.01.

4.2 Equilibrium selection

In this section we study the profile of equilibrium strategies with respect to the agents' initial beliefs.

We set the speed of learning process equal to 0.25 (that is $\delta = 0.25$), the number of iterations of GA per period equal to 25 (see Section 4.3 for more details) and consider 153 different compositions of initial population defined on the basis of agents' initial beliefs (any possible permutation of 16 agents for the possible three types of initial beliefs). For every possible composition of initial population we ran 50 simulations, modifying the seed of random numbers, so that we have $153 \times 50 = 7650$ observations.

The following Figure reports in a three dimensional simplex the frequency of the rest points of the simulations. Every vertex corresponds to a population composed by only one type of agents; in particular the vertex on the top corresponds to a population with only NA agents, the vertex on the right-bottom to a population of only WA agents and finally the vertex on the left-bottom to a population with only SA agents. Every point in the simplex corresponds to a mixture of these three types of agents and the darker is the color of that point the higher is the frequency with which the corresponding combination of types occurs in the simulation.

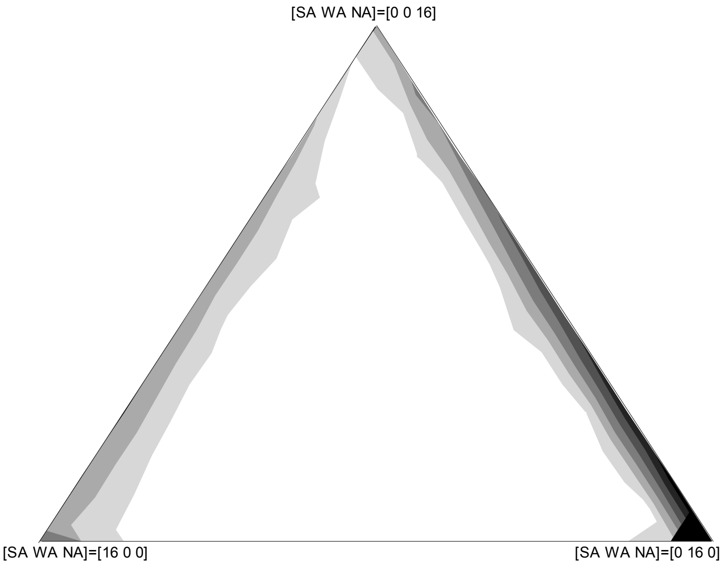

Figure 1

Figure 1 shows that most of the rest points is characterized by a population of WA and NA agents and only few simulations converge to a population with a significative amount of SA agents (in particular the shadow zone on the left represents more or less the 15% of total rest points, while the shadow zone on the right the 80%). It is worth observing that populations composed only by SA and NA agents (all the rest points near to the segment on the left) and those with only WA and NA agents (all the rest points near to the segment on the right) are the most frequent, while population composed by SA and WA agents show difficult to coexist (all the rest points near to the segment on the bottom).

However notice that Figure 1 does not show how the rest points are related to the initial composition of the population. To see this we calculate the number of times a simulation converged to every points in the grid (all the possible rest points) and consider only those for which this number is greater than 1/100 of the total number of simulations (that is $7560/100 = 75.6$); we find that most simulations converge to two main sets, one on the right, denoted as A, and one on the left, denote as B, of simplex. Moreover we characterize the basin of attraction of these two sets by calculating for every possible starting point the percentage of simulations converging to one of two sets. The following Figure reports the results:

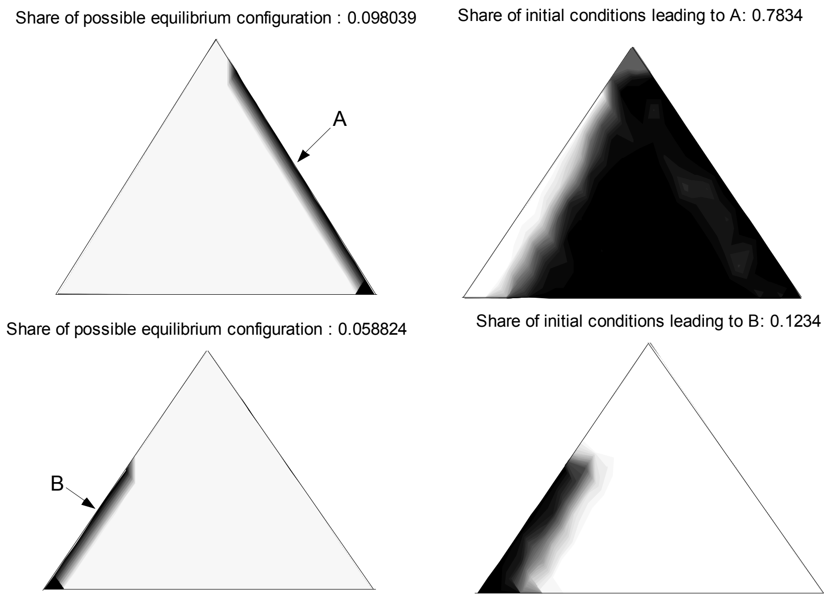

Figure 2

On the left of Figure 2 we represent the two attractor sets (A and B), while on the right the two basins of attraction, where the depth of the gray is proportional to the probability of convergence to the relative attractor set; we note that the set A is greater than set B (respectively 9.8% and 5.9% of all possible rest points), but the basin of attraction of the first one is greater than the second one (respectively 78.3% and 12.3% of the total number of simulations). Notice that the two attractor sets count for more than the 90% of the total number of simulations.

The light gray zone on the frontier of two basins represents the initial populations for which the convergence to one of two sets is not well defined, that is, starting from those configurations of the population, the probability of convergence to both sets is considerably greater than 0. For example, an initial population composed by 14 optimistic agents, 2 mildly optimistic agents and 0 pessimistic agents has a probability equal to 0.32 to converge to the set A and 0.56 to the set B, while an initial population composed by 12 optimistic agents, 2 mildly optimistic agents and 2 pessimistic agents has a probability equal to 0.50 to converge to the set A and 0.38 to the set B; heuristically the number of optimistic agents is crucial in deciding the probability of convergence to B.

Our previous findings are confirmed by the following Figure in which to every possible initial configuration of the population is associated the average payoff in the corresponding resting points:

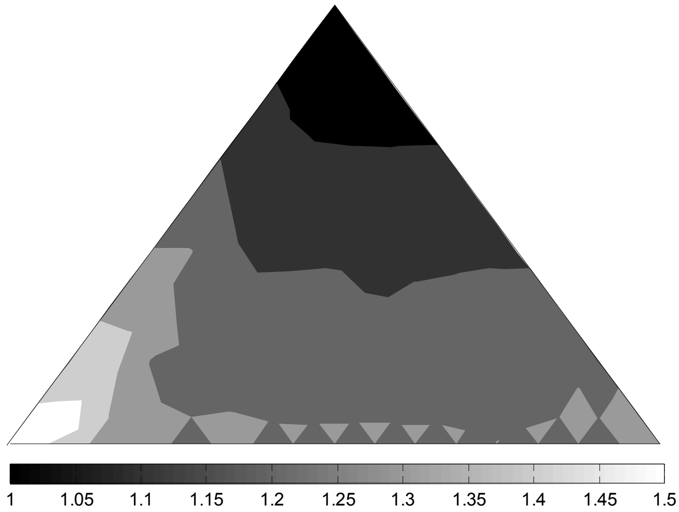

Figure 3

The low efficiency of the resting points is evident. The most efficient **CS**, corresponding to the CPNE, is given in our case by 5 coalition of size 3 and a coalition of size 1 and it corresponds to an average payoff of 1.5628. We see that only the simulations where almost all the population have optimistic initial beliefs converge to this result, otherwise the average payoff is notably lower. As we expected there is a strong relationship between the composition of initial population and the average payoff of the corresponding rest points. According to Figure 2 we see that if the number of mildly optimistic agents is greater than a certain threshold with respect to the number of optimistic agents, then the corresponding rest points are characterized by an average payoff lower than the one could be theoretically expected on the basis of the composition of the initial population (indeed the basin of attraction of the set A includes initial population where optimistic agents are a relevant part). This result is due to the free-riding behaviour of mildly optimistic agents in coalitions of size 3 (given their beliefs, this is their best action), which causes the optimistic agents to revise their beliefs downturn. However, if the number of mildly optimistic agents is very low, it is possible that the cooperative behavior of optimistic agents affect the beliefs of mildly optimistic agents, inducing the latter ones to cooperate in coalitions of size 3. The following Figure, reporting the difference between the effective average payoffs and the theoretical average payoffs calculated on the basis of the initial population, confirms this intuition:

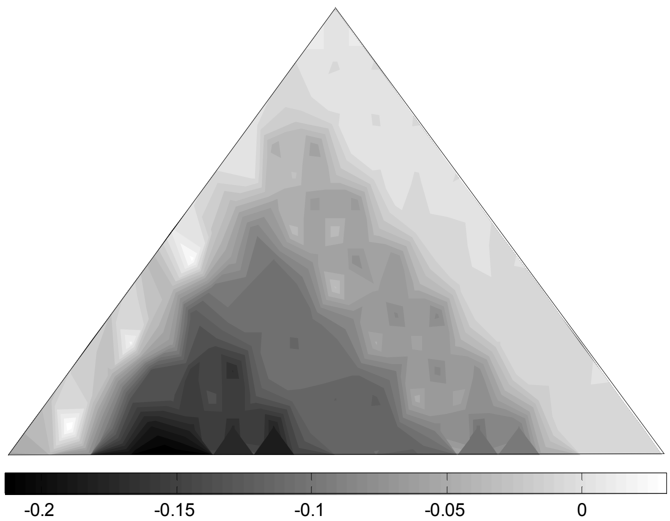

Figure 4

The initial populations where mildly optimistic agents are a few (1 or 2) and optimistic agents are the majority (from 8 to 16) converge towards the set B and therefore the effective average payoff is higher than the theoretical one (in Figure 3 this is represented by the three small light areas on the left); in other words the mildly optimistic agents become optimistic agents and cooperate in coalitions of size 3 (this is also evident from Figure 2 if we consider the extent of the basin of attraction of the set B). On the contrary, when the number of mildly optimistic agents is relevant (from 3 to 8), even if optimistic agents are always a significant part of the initial population (from 13 to 7), the system converges with higher probability to the set A (see Figure 2) and therefore the effective average payoffs are considerably lower than the theoretical ones.

4.3 Agents' computational ability and learning speed

There are important contributions in the literature showing the importance of agents' computational abilities in the selection of an equilibrium (see [1]), as well as of their learning speed (see [14]). In our context the computational ability is measured by the number of iterations of GA per period, while the speed of learning is measured by δ. In the following we analyze how our findings are affected by these two parameters and provide example of simulations with alternative values.

4.3.1 Computational ability

It is evident that agents' computational abilities can affect the rest points of the simulation, since they affect the accuracy of the individual maximization process. Therefore it is worth to analyze this aspect in order to eliminate a possible bias in the results. We ran several simulations with 16 optimistic agents, setting $\delta = 0.25$ and varying the number of iterations of GA performed in every period from 1 to 50. The following Figure reports the most significative statistics (average coalition cardinality, average number of cooperators, average payoffs, average convergence time with respect to the number of iterations of GA for period)[16]:

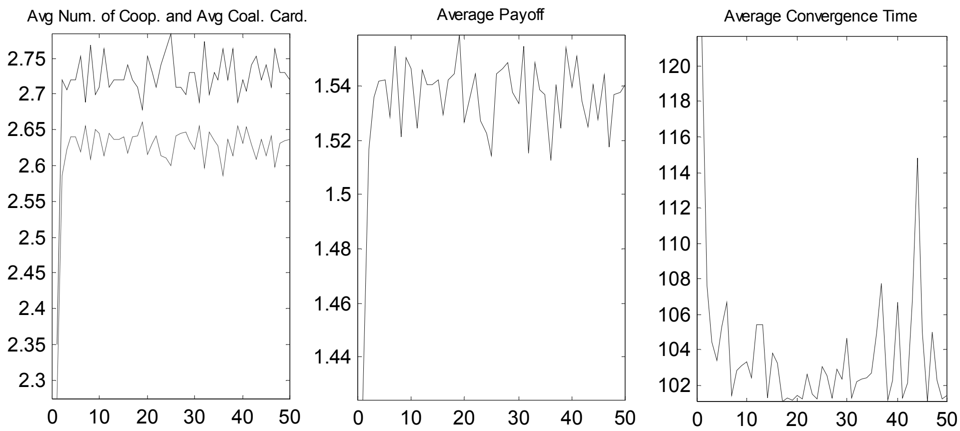

Figure 5

Figure 5 shows that a number of iterations of GA equal to 25, that is the one we have considered in the previous Section, guarantees that the rest points are not substantially affected by the agents' computational ability (indeed this already holds for a number equal to 5). We note that for a low number of iterations of GA, that is in presence of "imprecisions" in agents' decision process, both the average cardinality and the average payoff are remarkably lower.

To investigate thoroughly this aspect we ran a simulation where the number of iterations of GA is set to 1. The following Figure reports the results:

[16] The top line in the first picture of Figure 5 depicts the average cardinality of the formed coalitions, while the lower line reports the average number of cooperators when the number of iterations of GA varies as reported on the horizontal axis.

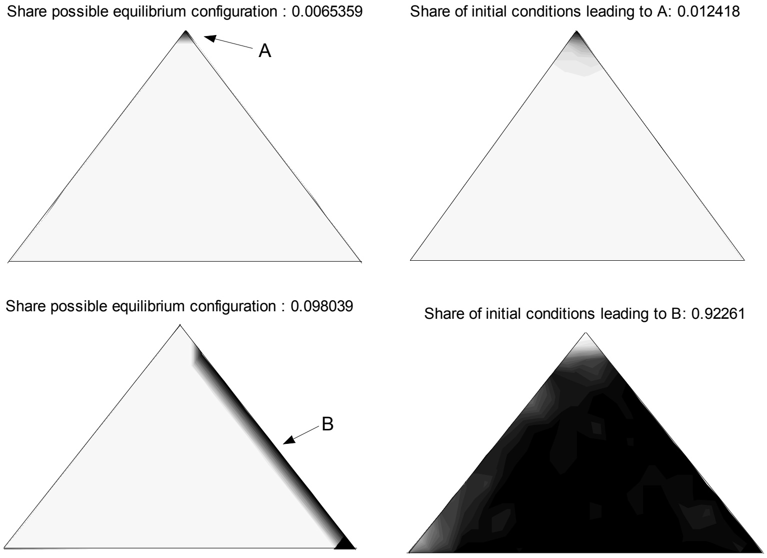

Figure 6

As we expected, the greater imprecision in the decision process leads to the formation of coalitions of smaller cardinality; in particular the attractor set on the left disappears, the one on the right collects more than the 92% of the total number of rest points and another attractor set appears, whose size, however, is very small. The implications in terms of overall average payoffs are trivial, so that it is redundant to report the Figure.

This findings can be heuristically explained by the fact that there are many situations where it is sufficient the presence of just one defector within a coalition to make defection the best action the next period for all the agents belonging to that coalition and the inaccuracy in the individual reasoning may induce someone to choose the wrong strategy, i.e. one prescribing defection when cooperation would have been profitable. This mistake will certainly lower current payoff but, more importantly, will propagate the incentive to defect in all coalitions of that size, due to the working of the expectation revision mechanism.

4.3.2 Learning speed

Another key factor in evolutionary games is agents' learning speed. In order to test the importance of this factor we ran several experiments varying the parameter δ, taken as given a population of 16 optimistic agents and setting to 25 the number of iterations of GA for period; in particular we consider 101 different values of δ in the range $[0, 1]$ (that is from 0 to 1 with step 0.01) and for any value of δ we ran 50 simulations. The following Figure reports the most significative statistics, whose meaning is the same as in Figure 5:

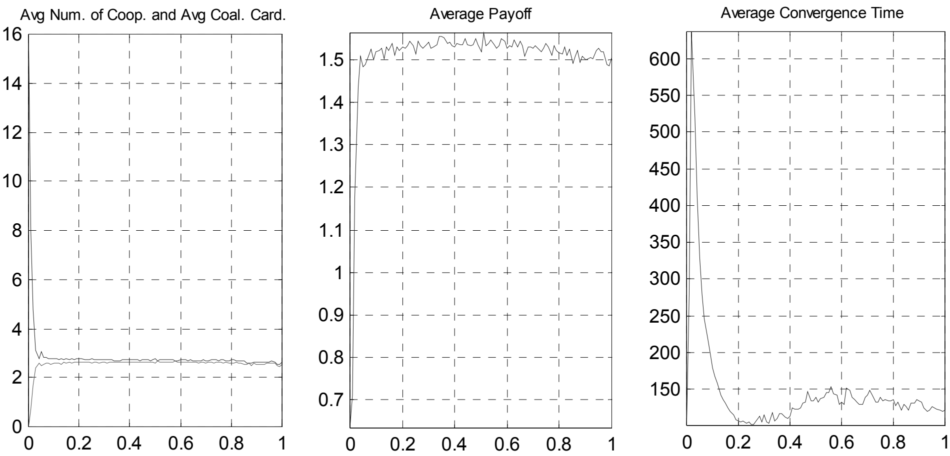

Figure 7

We observe that for $\delta \in [0.1, 0.9]$ the properties of rest points do not show substantial changes, so that the choice of $\delta = 0.25$ for our previous experiments seems to be appropriate; however it is worth to investigate what happens for extreme values of δ.

For low value of δ agents do not learn from experience; they play strategies indicating an acceptable coalitional cardinality greater than that compatible with full cooperation and, at the same time, they do not cooperate, given the persistent (incorrect) beliefs that other agents will be cooperating. As learning speed increases, agents learn that this strategy is also adopted by other agents, so that they decrease the acceptable coalition cardinality and play cooperation (this happens for values of δ around 0.1). The following Figure reports the attractor sets and the basins of attraction for $\delta = 0.05$:

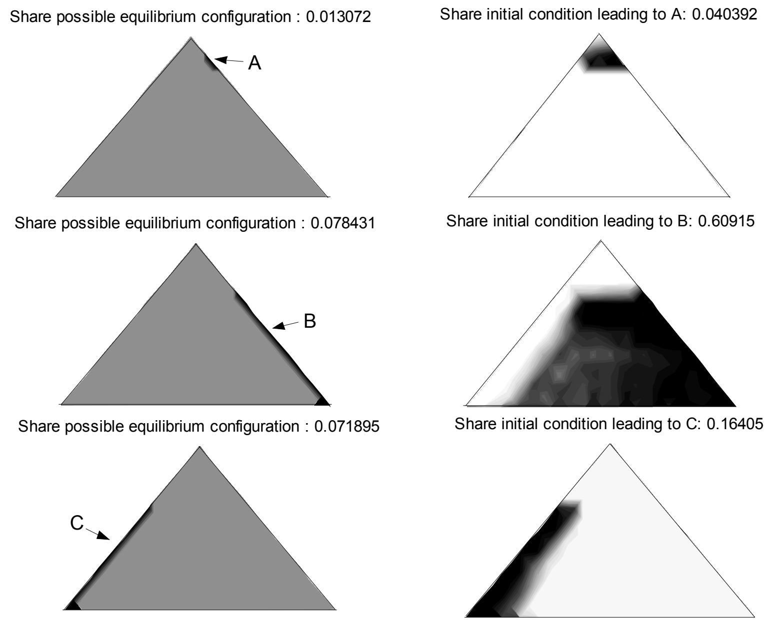

Figure 8

With respect to the case in which $\delta = 0.25$, there is a new attractor set, denoted as A in Figure 8, characterized by a population with a large number of NA agents; however there is a sensible reduction in the size of the basin of attraction of the set B. These results confirm our previous intuition, even if the basin of attraction of the set C, representing the most efficient rest points, is slightly increased.

Finally for high values of δ we notice that the average payoff is slightly decreased. This happens because a higher "reaction" to experience can easily lead to the destruction of optimistic beliefs in presence of possible "mistakes" in the formulation of strategies. For example if some agent does not cooperate in a coalition with cardinality equal or lower than 3 (we know that this is not the best action for an optimistic agent), this can lead other optimistic agents to revise drastically their beliefs, so that they will not cooperate in future in coalitions of cardinality equal to 3. This effect is stronger when the value of δ increases. The following Figures reports the attractor sets and the basins of attraction for $\delta = 0.95$:

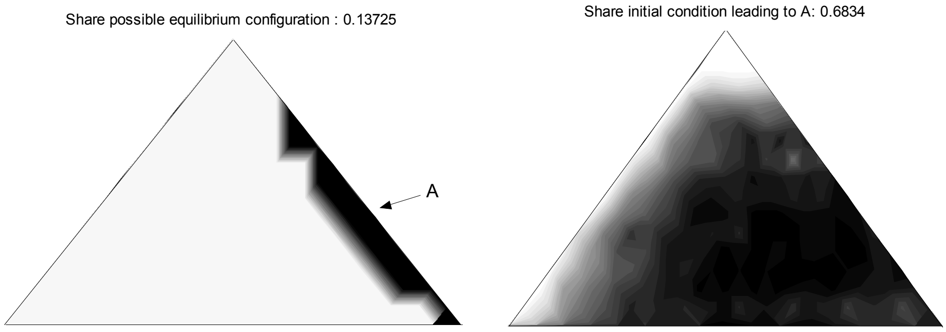

Figure 9

Notice that, with respect to the case in which $\delta = 0.25$, the attractor set with SA agents disappears, which confirms our previous intuition (of course this result entails a decrease of the overall efficiency).

5 Conclusions

In this paper we have shown that, if agents adjust their behaviour on the basis of their experience, then they play the CPNE strategies only if the population starts out with a priori beliefs that are close to those sustaining coalition-proof outcomes; in this case the role of the learning process will be just to size down the dimension of the equilibrium coalitions to the one consistent with individual incentives to cooperate. The composition of initial population strongly affects the selection of equilibrium. The computational experiments highlight how there are two main attractor sets, the first one characterized by a population composed by a large part of strongly associative agents and a minority of not associative agents, while the second one is characterized by a population composed by a large part of weakly associative agents and a minority of not associative agents. This suggests that equilibria with a population of strongly associative and weakly associative agents are not "sustainable" because of the free-riding behaviour in large coalitions of the weakly associative agents. Moreover we observe that the possibility that weakly associative agents become strongly associative is limited to the case in which weakly associative agents are a strict minority in a population with a large share of associative agents. On the contrary, if the number of weakly associative agents is sufficiently high, we observe convergence toward an equilibrium in which there are no strongly associative agents. These findings suggest that these two attractor sets can be preserved even in face of small random perturbations in agents' beliefs. Therefore a straightforward extension of this work would be to allow for random mutations in agents characteristics in order to test the robustness of the attractor sets from a purely evolutionary point of view.

Finally we test the robustness of our results for various parametrizations; we find that the overall efficiency is an increasing function of agent's computational capabilities, while an increase in agents' learning speed can have an ambiguous effect.

A Geometric selection

Let F_i^t be the vector representing the fitness at the period t of the strings of the agents i. The geometric selection assigns to the string j the following probability to be selected

$$P_j = \left[\frac{p^{gs}}{1-(1-p^{gs})^J}\right](1-p^{gs})^{r_j-1},$$

where p^{gs} is the probability of the best string to be selected, J the total number of strings and r_j the ranking of the string j, with the highest fitness having r equal to 1.

B Stopping condition

The stopping condition is implemented taking in account the possibility that a sequence of coalition structures can show a cyclical pattern or that possible temporary departures from steady state can occur just because of the intrinsic stochasticity of agents' decision process. Therefore we adopt a stopping condition which tests if a change in the coalition structure can be explained by simple random disturbance around a structural attractor. In particular, in period t, we consider the distribution of coalitions-cooperators in the periods from $t-n$ to $t-1$ (we choose $n=50$ in our computational experiments) and verify that the current period observation is drawn from the same distribution generating the past observations. In particular, let O^t be the $I \times (I+1)$ matrix in which the number at the crossing of row r and column c reports the number of agents that are allocated in a coalition of cardinality r with c cooperators, and let $T^t = \frac{1}{n} \cdot \sum_{s=t-n}^{t-1} O^s$; then we calculate the following index[17]

$$\chi = \sum_{i=1}^{I}\sum_{j=1}^{I+1}\frac{\left(O_{ij}^t - T_{ij}^t\right)^2}{T_{ij}^t}$$

and the simulation is stopped when χ is lower than the value of the χ^2 statistics at level 0.99 with $((I-1)\cdot I)$ degrees of freedom.

[17] If $h_{i,j}^t$ is zero, then in the computing the test we consider a small positive number.

References

[1] Anderlini, L. and H. Sabourian (1995). "Cooperation and Effective Computability". *Econometrica*, 63(6), 1337-1369.

[2] Arifovic, J. (1994). "Genetic Algorithm Learning and the Cobweb Model". *Journal of Economic Dynamics and Control*, 18(1), 3-28.

[3] Arifovic, J. (1995). "Genetic Algorithms and Inflationary Economics" *Journal of Monetary Economics*, 36, 219-243.

[4] Arifovic, J. (1996). "The Behaviour of the Exchange Rate in the Genetic Algorithm and Experimental Economies". *Journal of Political Economy*, 104(3), 510-541.

[5] Arthur, W. B. (1991). "Designing Economic Agents that Act Like Human Agents: A Behavioral Approach to Bounded Rationality". *American Economic Review*, 81(2), 353-59.

[6] Axelrod, R. (1992). "The Evolution of Strategies in the Iterated Prisoner's Dilemma". Mimeo.

[7] Bernheim, B. D., B. Peleg and M. D. Whinston (1987). "Coalition-Proof Nash Equilibria: Concepts". *Journal of Economic Theory*, 42(1), 1-12.

[8] Binmore, K. (1988). "Modeling Rational Players: Part II". *Economics and Philosophy*, 4(1), 9-55.

[9] Birchenhall, C.R. (1996). "Evolutionary Games and Genetic Algorithms". In *Computational Economic Systems.* Edited by Gillì M., pp. 3-23. Kluwer Academic.

[10] Bullard, J. and J. Duffy (1998). " A model of learning and emulation with artificial adaptive agents". *Journal of Economic Dynamics and Control*, 22, pag. 179-207.

[11] de Vany, A. (1996). "The Emergence and Evolution of Self-Organized Coalitions". In *Computational Economic Systems.* Edited by Gillì M., pp. 25-50. Kluwer Academic.

[12] Farrell, J. and S. Scotchmer (1988). "Partnerships". *Quarterly Journal of Economics*, 103(2), 279-97.

[13] Fiaschi, D. and P.M. Pacini (1998). “Endogenous Coalition Formation with Identical Agents”. *Working paper n° 308* (Progetto d’Ateneo: L’economia italiana e la sua collocazione internazionale: una redifinizione delle politiche di welfare e dell’occupazione per una più efficiente crescita economica), Dipartimento di Scienze Economiche, Università degli Studi di Bologna.

[14] Gale, J., K. G. Binmore and L. Samuelson (1995). ”Learning to be Imperfect: the Ultimatum Game”. *Games and Economic Behavior*, 8, pp.56-90.

[15] Greenberg, J. (1990). *The Theory of Social Institutions. An Alternative Game-Theoretic Approach.* Cambridge University Press, Cambridge, UK.

[16] Greenberg, J. (1994). “Coalition Structures”. In *Handbook of Game Theory.* Edited by Aumann, R. and S. Hart, vol. 2, chap. 37, pp. 1305-37. North-Holland, Amsterdam.

[17] Guesnerie, R. and C. Oddou (1988). “Increasing Returns to Size and Their Limits.” *Scandinavian Journal of Economics*, 90, 259-73.

[18] Herrenstein, R. (1991). “Experiments on Stable Suboptimality in Individual Behavior”. *American Economic Review, Papers and Proceeding*, 81(2), 360-64.

[19] Holland, J.H., K.J. Holyoak, R.N. and P. Thagard (1986). *Induction: Processes of Inference, Learning and Discovery.* MIT Press, Cambridge Mass.

[20] Holland, J.H. and J.H. Miller (1991). “Artificial Adaptive Agents in Economic Theory”. *American Economic Review, Papers and Proceeding*, 81(2), 365-70.

[21] Mailath, G. J. (1998). ”Do People Play Nash Equilibrium? Lessons From Evolutionary Game Theory”. *Journal of Economic Literature*, Vol. 36, September, 1347-1374.

[22] Marimon, R., E.M. McGrattan and T.J. Sargent (1990). “Money as a Medium of Exchange in an Economy with Artificiallly Intelligent Agents”. *Journal of Economic Dynamics and Control*, 14, 329-373.

[23] Marimon, R. (1996). “Learning from Learning in Economics”. Working Paper N. 96/12, European University Institute, Florence.

[24] Miller, John H. (1996)a. "Evolving Information Processing Organizations ". Mimeo.

[25] Miller, John H. (1996)b. "The Coevolution of Automata in the Repeated Prisoner's Dilemma". *Journal of Economic Behavior and Organization*, 29(1), 87-112.

[26] Young, H. Peyton and Dean Foster (1991). "Cooperation in the Short and in the Long Run". *Games and Economic Behavior*, 3(1), 1991.