Agent Compatibility and Coalition Formation: Investigating Two Interacting Negotiation Strategies

Carlos Merida-Campos and Steven Willmott

Universitat Politècnica de Catalunya, Software Department, E-08034 Barcelona, Spain
{dmerida,steve}@lsi.upc.edu

Abstract. This paper focuses on the Coalition Formation paradigm as a market mechanism. Concretely, Coalition Formation occurs as part of a wider open world and may occur many times during the lifetime of a population of agents. This fact can in some circumstances be exploited by agents to re-use existing partial coalition and social relationships over time to improve Coalition Formation efficiency. The aim of the work is to analyze the dynamics of two concrete rational behaviors (Competitive and Conservative strategies) and, in particular, to investigate how agents in a heterogeneous population cluster together across multiple Coalition Formation episodes and varying tasks. Preliminary resuls are also shown regarding the manner in which playing distinct strategies interact with one another.

1 Introduction

One branch of economics research investigates how specific market situations force certain behaviors in their participants. A common example of this phenomenon is *Smith's Invisible Hand*, under which, in competitive market scenarios, traders are forced to set the prices where the supply meets the demand [13]. This specific situation does not appear too often in its pure form, and traders must often choose in what could be a large space of different pricing strategies.

The majority of agent-based research in the Economics area, focuses on studying strategy interactions and dynamics given a certain market. In this paper we focus on the related question of strategies and dynamics in market mechanism scenarios where Coalition Formation is required. Coalition Formation is a problem which has been extensively studied from Multi-Agent Systems perspective. The problem centers on finding subsets of agents from a general population to form groups which can most effectively carry out a particular task. The area includes research on a variety of different fronts, including problems such as finding core allocations of payoff between the members of a coalition [1], finding stability properties between coalition structures [7], finding solutions to a problem in a cooperative way [5], or finding solutions to a coalition problem in a competitive way, where coalitions compete amongst themselves for a payoff [12].

M. Fasli and O. Shehory (Eds.): TADA/AMEC 2006, LNAI 4452, pp. 75–89, 2007.

In the market mechanism we suggest here, coalition formation occurs as part of a wider open world and may occur many times during the lifetime of a population of agents. As shown in [12] this fact can in some circumstances be exploited by agents to re-use partial coalition and social relationships over time to improve Coalition Formation efficiency. Such a broader perspective however raises interesting questions for coalition formation environments. This paper goes further in analyzing the dynamics of two concrete rational behaviors, a Competitive Strategy and a Conservative Strategy within such an environment. These strategies are common examples of rational behavior in economic markets. Whilst there are more complex strategies, these two already show significant interesting behavior.

This type of continuously running market environment requiring coalitions reflects a large number of real world scenarios such as bidding for construction projects, the formation of consortia for product development or strategic alliances in emerging markets. In this context we aim to answer the following questions:

– To what extent do specific strategies affect the type of coalitions formed?
– To what extent do specific strategies affect the stability properties of the system?
– How do different agents with different strategies affect each other when they interact?
– Which strategies benefit agents and the general population more in which situations?

The paper is structured as follows: Section 2 explains the concrete market mechanism we are using, the Iterative RFP Coalition Formation method as well as examples to illustrate the method. Section 3 explains the agent based system designed to model the Iterative RFP Coalition Formation Method. Section 4 explains the characteristics of the strategies we are testing in this paper, and analyzes them from a theoretical perspective. Section 5 explains the experimental setup and the results obtained to underpin the theoretical analysis and to provide some new information. Section 6 covers related work and finally, Sections 7 and 8 provide conclusions and ongoing / future work repsecitvely.

2 Iterative RFP Coalition Formation Method

As the basis for experimentation, analysis in this paper adopts a model of worlds based on Request For Proposal (RFP from now on) scenarios. This model was first studied by [11], and further explored in [12]. In this environment, an entity or entities regularly issues a call for tender to provide specific goods or services with certain characteristics. Providers compete amongst themselves (either individually or in cosortia – *coalitions*). Providers and/or coalitions bidding for a particular call are ranked according to an evaluation of their skills for the task, and receive a payoff according to their placement in the ranking.

There are many existent real systems that follow the RFP type procedures such as public building projects, competitive tender for government contracts or even collaborative research project grants. RFP environments can also be seen as emerging market opportunities in an economy, with individual calls for tender representing new opportunities for profit. Such system are characterized as follows:

– Agents, or groups of agents compete for a given goal.
– The best agents, or groups of agents are rewarded with the award of the contract to carry out the task and its subsequently payoff.
– Some systems, may also have policies for rewarding 2nd, 3rd, 4th etc. ranked agents / groups of agents.
– The process repeats over time.
– The objective agents are competing for varies over time – with different CFPs issued corresponding to functionally different tasks. In this way, an agent or a group of agents that were very competent for a certain goal, could become weak for a different one.
– Groups could be dynamic and might change depending on the market situation.
– Agents are individual utility maximizers. And share the same preference, that is payoff maximization, but might have different strategies for maximizing their utility.

A protocol with the described characteristics could also potentially be used as a service composition mechanism in an automated environment where agents representing services team up to compete with other teams to present the most competitive package of services given certain requirements.

3 Problem Definition and Agent Based Models

For the purposes of this paper, the problem is formalized in the following way: In every game g there is a task T_g that is defined by a set of K tuples. Each tuple represents an skill and its corresponding demanding value for the named task:

$$T_g = \{\langle sk_0, T_{g0}\rangle, \langle sk_1, T_{g1}\rangle, \ldots, \langle sk_K, T_{gK}\rangle\}$$

Every agent A_i in the population has a certain expertise degree in each of the K skills that task T_g is defined with:

$$A_i = \{\langle sk_0, A_{i0}\rangle, \langle sk_1, A_{i1}\rangle, \ldots, \langle sk_K, A_{iK}\rangle\}$$

A coalition C_j is a a set of one or more agents $\{A_x, .., A_z\}$. Agents' skills are aggregated in the coalition in such a way that the resultant values of that aggregation represents the skills of the coalition entity. This is noted as:

$$C_j = \{\langle sk_0, C_{j0}\rangle, \langle sk_1, C_{j1}\rangle, \ldots, \langle sk_K, C_{jK}\rangle\}$$

The concrete aggregation function used is:

$$C_{jp} = \max_{\forall q: A_q \in C_j} A_{qp} \qquad \text{Having } 0 \leq p \leq K. \tag{1}$$

We could consider each skill as a necessary subtask for performing task T_g. In this way, by using the aggregation function shown in equation 1, the agent in a coalition which is the best fit for performing a certain subtask will be in charge of it.

Amongst the different possibilities for aggregating agent skills in a coalition, function 1 has been choosen as is a reasonable metaphor of many real coalitional processes. For example, if we consider a consortium of partners participating in a call for proposals, each member of the consortium will be representative of a certain part of the proposal, and normally is the partner best fitted for that part of the work.

Coalition C_j is endowed with a certain score $scr(C_j, T_g)$. This score is negative if the coalition is *non competent* in all the skills for performing the task ($\exists p : (0 \leq C_{jp} < T_{gp})$), and is positive otherwise. More concretely, the functions used for the case of existence of *non competent* skills in the coalition is:

$$scr(C_j, T_g) = -\#sk_p : (0 \leq C_{jp} < T_{gp}) \tag{2}$$

When Coalition C_j is competent in every skill ($\forall p : (0 < T_{gp} \leq C_{jp})$), the function used is:

$$scr(C_j, T_g) = \sum_{p=0}^{K} C_{jp} - T_{gp} \tag{3}$$

As we can see in Equation 2, the score of a coalition with some non competent skill (*non competent coalition*), is the negative value of the number of non competent skills. In this way, the more skills in which the coalition is not competent, the lower its score will be.

In Equation 3, we can see that the score of a coalition, competent in every skill (*competent coalition*) is the sum of excess value in every requested skill. In both equations, we can see that only those skills with value higher than 0 count for their evaluation. Those with values equal to 0 are ignored. This represents the fact that some tasks do not need a certain skill to be performed, and so the degree of ability of a coalition in that skill is not taken into account for its evaluation.

There is non-linear mapping from coalition score to coalition payoff, as in our model, the payoff of a coalition does not depend only on its score but also on the scores of other coalitions. All coalitions are decreasingly ordered by score, and are priced according to their rank with an exponentially decreasing amount from the best one to the worst. Agents within a coalition spread the coalition payoff evenly. The concrete payout function we use is:

$$\text{pay(rank)} = \begin{cases} \text{MaxAmount}/2^{\text{rank}-2}, & \text{for the last competent coalition} \\ \text{MaxAmount}/2^{\text{rank}-1}, & \text{for the other competent coalitions} \end{cases} \tag{4}$$

Amongst the different possibilities for doing the mapping, this concrete exponential function has been chosen for two reasons: first, because this function has the property that, independently of the number of competent coalitions, the total amount of money spread will always be $MaxAmount$. This is an interesting property in order to compare the economic behavior of different populations. The second reason is because this represents an exponential distribution of wealth for which there is empirical evidence from Real Economies (see [6]). This function creates a rich set of possible scenarios in which coalitions optimize their trade-offs between score and size, growing in size only when it is valuable to do so.

For every game, a subset of agents are asked at random about an action to take towards its membership in a coalition. The choices that an agent has are:

– Stay in the coalition.
– Stay in the coalition optimizing it by firing (expelling) one or more members.
– Leave the coalition in order to join a different one.
– Leave the coalition in order to replace one or more agents in a different one.
– Create a new coalition.

When an agent's decision involve firing or replacing an agent, this action is submitted to the coalition who will evaluate it. In order to get the action accepted and executed, it must be approved for more than the half of the members of the coalition affected, otherwise the action is rejected and not performed. In that case that an agent is fired or replaced, it automatically becomes the only member of a brand new coalition.

When agents are requested to perform an action they are able to submit as many proposals as they want, if none of them is accepted the agent remains in the same coalition.

Agents are farsighted in the sense that they know the score and payoff values of any action the agent wants to consider prior to its submission.

4 Conservative and Competitive Strategies

This paper shows results on the different outcomes obtained by using two different strategies: *competitive* and *conservative*. Agents who choose a *conservative* strategy make decisions on which coalition to join based on the payoff this coalition is expected to have. Agents who choose a *competitive* strategy make decisions on which coalition to join based on the score that this coalition is expected to have. Both strategies are myopically rational, as when an agent using some of those strategies is asked to make a choice, it counts with all the information available at that instant of time, concretely, they have the information on potential payoff and score of any coalition they could create by carrying out any of the 5 possible actions defined in the previous section, but they do not count on the possible reactions of the other agents after its decision has been performed.

Both strategies are arguably rational. For the case of *conservative* strategy, it makes sense to choose the most profitable coalition at a certain time, expecting that the situation will not change until the end of the game. For the case of *competitive* strategy, it makes sense to join the coalition with highest score, as even if the payoff is worse than in another coalition with lower score but fewer members to share the benefit, a new member could be attracted by this growing score coalition and make it grow in the ranking and gain a higher payoff. In other words, for the case of *conservative* strategy, the agent expects to get the maximum payoff at any time, while for *competitive* strategy, the agent invests for a future better payoff as a side effect of being in a highly competent coalition.

4.1 Theoretical Dynamics of Competitive Population

Competitive agents try to be in a coalition with the highest possible score. At the same time, coalitions of Competitive agents accept joining proposals, optimizations or replacements proposed as long as they improve the coalition score.[1] This behavior implies that, for a given task, each movement of an agent from/to a coalition of Competitive agents, involves an improvement in its score, otherwise the agent would not had had any motivation to move from/to there. Thus the score of the best coalition in a game for a given task is monotonically increasing.

The maximum score of a coalition is obtained with a coalition of as many members as skills required for the task (at maximum). Agents do not create coalitions with more members than skills, as given aggregation function 1, there is just one agent providing the maximum value for each skill, then if the coalition has more agents than required skills, it could be optimised by expelling those agents who do not provide any maximum value to any skill. Due to this, Competitive agents in the RFP model do not create the grand coalition, hence this strategy configures a non-superadditive game.

As the coalition size is self-limited by a certain maximum size, and the score of the best coalition is monotonically increasing, the score of the best coalition will stabilise at a certain point while the task doesn't change. The same happens with the score of the second best coalition, and so on. This way, a system with a population of Competitive agents competing for a given task that do not change, converges into an stable state.

4.2 Theoretical Dynamics of Pure Conservative Population

Conservative agents try to be in a coalition with the highest possible per agent payoff. At the same time, coalitions of Conservative agents accept joining proposals, optimizations or replacements proposed as long as they improve the coalition per agent score.[2] This behavior ensures that during the process of coalition formation, each movement of an agent from/to a coalition of Conservative agents, involves an improvement in the per agent payoff of the coalition otherwise, the agent would not have had any motivation to move from/to there. In order to improve the per-agent payoff, a coalition can either improve its score to increase its ranking, or it can reduce its size having less members to split the gains between less members. This second way of improving payoff implies in some cases a reduction of coalition score.

[1] As a secondary criterium of acceptance, we use the size of the coalition, i.e. if the coalition score is the same, a competitive agent will propose/accept when the number of agents is smaller than in the origin/original coalition.

[2] The same secondary criterion of acceptance as in competitive strategy is applied.

Differently to the best coalition's score, the best coalition's payoff is therefore not monotonically increasing. The best individual payment can decrease when the coalition that has it, is out-ranked by another coalition. In this case, the previous best payed coalition receives less payment as it is in a lower rank, and the out-ranking coalition, might have more members than the previous best-paying coalition, thus lower individual payoff. Under these conditions of non monotonicity, convergence to an stable state cannot be ensured in an environment with Conservative agents.

The sizes of coalitions are determined by the concrete skill distribution amongst the population, the requirements of the task and the payoff function. More concretely, the growing possibilities are determined by score differences from coalitions in the ranking. Let A_1 be an agent, C_1 be a coalition and $\backslash C_1$ the rest of coalitions competing at a certain moment. The ranking of C_1 is noted as $rank(C_1, \backslash C_1)$, its size as $|C_1|$, and its payoff as: $pay(rank(C_1, \backslash C_1))/|C_1|$. We can claim that C_1 will never grow in size as long as:

$$\nexists A_1 : pay(rank(C_1 \cup A_1))/(|C_1| + 1) > pay(rank(C_1))/|C_1| \tag{5}$$

Then the difficulty of a coalition has in growing depends on payoff function. Since the payoff function we are using in our model is monotonically decreasing as a function of the ranking of the coalition, in order to have a higher payoff in coalition $C_1 \cup A_1$, the following must be true:

$$rank(C_1 \cup A_1) > rank(C_1) \tag{6}$$

From this, we can see that the possibilities of a coalition to grow up in size also depend on how difficult what is stated in Condition 6 is.[3] As a matter of fact, the difficulty of condition 6 to hold for any arising coalitions, depends on the skill distribution of agents, and on the size of coalitions, in such a way that the smaller the coalition, the easier it is for condition 6 to be fulfilled. The Influence of coalition size is explained as follows: in a coalition with no redundant agents, each agent is giving the maximum value to one or more skills. If the coalition has few agents, its members will each have more skills to be responsible for. If agents of a population have the same total sum of skill values $(\sum_{p=0}^{K}(A_{ip}))$, there will be many agents with different skill values in the distribution that could improve the coalition score by taking responsibility of many more skills than if the coalition were large and each skill would be responsibility of a certain "expert" agent. In this way, by having more chances of increasing coalition score by raising the value of more skills, we will have also more chances of rising the coalition rank.

The stability analysis in this population is more complex than in competitive population. The system will be stable once agents have reached Pareto optimality, and Nash equilibrium. This situation is made more complex by different

[3] Note that it could happen that joining more than one agent to the coalition $(C_1 \cup A_1 \cup \ldots A_n)$ we could improve the coalition enough to make it profitable for all the agents in it, but given our model, agents do not take decisions in a coordinated manner.

factors. An important one is that if coalitions do not need to grow too much to be competent, we will have many small coalitions, and as we have seen in the analysis of the importance of the coalition size, for the case of small coalitions it will be easier for agents to find profitable coalitions outside their current coalition. This could create a continuous movement of agents from coalition to coalition.

5 Experiments

To examine the dynamics of the strategies explained and validate the theoretical analysis, a range of experiments were conducted by simulation. Each simulation run consists of a set of a fixed number of model iterations, where agents follow the Iterated *RFP* protocol explained in section 2, for solving a number of different tasks that change sequentially after a fixed number of games.

In order to visualize the relationships established between agents in the experiments, we used Pajek [3]. Graph figures represent the relationships created between agents throughout a series of games of one or more experiments. Each node is an agent, and a link represents a collaboration that existed when the coalition was evaluated in a certain game. Agents collaborate when they are together in a competent coalition. In this way, a coalition is represented by a clique of connections amongst the agents in the coalition. In order to ease the visual analysis edges are colored depending on the frequency of the relationship; the more often a collaboration happens, the darker the line appears. The graph is represented using *Kamada-Kawai* algorithm implemented in *Pajek* that places nodes in a close position when they are connected with links of relative high value. In our case, agents appear close to each other when they have had frequent relationships. Throughout the rest of the paper, these graphs are named *collaboration graphs*.

5.1 Experimental Set-Up

A variety of configurations and parameterizations have been used in order to be able to check the statistical validity of outcomes. The underlying configuration for the experiments was:

- A static population of 100 agents, with abilities randomly distributed across 10 different skills. Each skill of an agent is assigned a positive integer score or zero. Agents are each assigned 200 skill points randomly distributed across their skills.
- A set of 100 different Tasks. Each one requiring a total of 100 skill points distributed randomly across the same 10 different skills.
- Each task is issued during 1000 games. And each game picks 25 agents at random to make a choice on their coalition preference (and hence potentially adapt their coalition according to the schemes defined in the previous section).

– Every competent coalition is subsequently ranked and rewarded according to function 4. The concrete *MaxAmount* value used for this function is 100. This represents the total amount that will be spread in each game amongst the competent coalitions.

Collaboration Graphs are used to monitor each relationship established at the end of every round. The payoff data for each one of the agents, as well as the coalition sizes are also monitored and analysed.

5.2 Competitive Behavior Experiments

A population of 100 agents following the Competitive behavior strategy have been used. Figure 1(a) shows the *collaboration graph* of the population in just one experiment, and figure 1(b) represents the aggregated results amongst the set of 100 experiments with different tasks to fulfil. In the first figure we can appreciate 30 vertices out of the existing 100, this is because the algorithm places vertices close to each other when relationships between them are very strong (in our case, strength values means frequent relationships), that means that some coalitions of agents repeat very frequently over time. In the same figure, a few residual low frequency relationships of some agents with some of other clusters can be appreciated. Those two facts verify the analysis performed in section 4.1. Concretely, the experiments verify that a population of agents using Competitive strategy converges to a stable state where at a certain point of the process, there is no movement of agents between coalitions. Those clusters of nodes in the graph that repeat very frequently are the structures created after the optimization process has finished. Residual relationships (relationship with low frequency) belong to the period in which the experiment is not stable yet.

The experiments also reveal another fact that validates the theoretical analysis. The size of coalitions tends to be equal to the number of requested skills by the task, having normally an expert agent on charge of every skill.

The second figure shows us that even when different tasks are used in the same experiment, the aggregated collaboration graph still reveals clear clusters of frequent collaboration. This is explained by the following fact: Competitive agents create coalitions that maximize the value of all the skills requested by a task. In order to do that, they create coalitions as big as necessary, having an expert agent for each requested skill, i.e. an agent which only contribution is to have the biggest value in the coalition for an specific skill. The only difference in coalitions created for different tasks is in the exclusion or inclusion of expert agents for skills that in some tasks are requested and in others are not. This way a substantial part of a coalition remains static from task to task, as in average, tasks issued usually have no more than 2 non-required skills (that vary from task to task) out of 10 skill.

5.3 Conservative Behavior Experiments

In these experiments, we use a population of 100 agents following the conservative behavior strategy. Figure 1(c) shows the *collaboration graph* of the

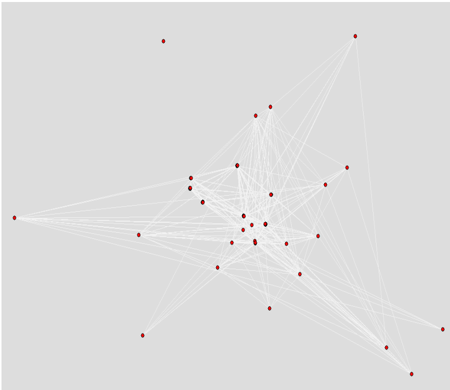

(a) 1 Experiment Result pure competitive population.

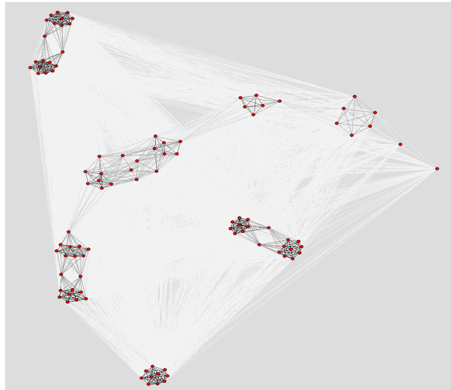

(b) 100 Experiments Result pure competitive population.

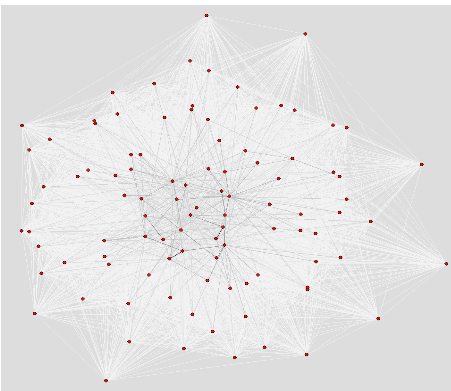

(c) 1 Experiment Result pure conservative population.

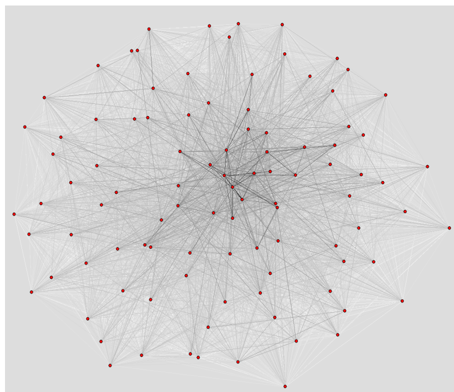

(d) 100 Experiments Results pure conservative population.

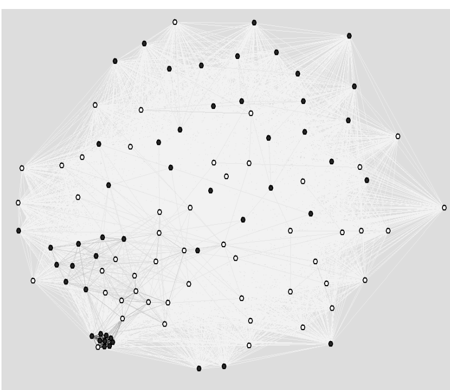

(e) 1 Experiment Result mixed population.

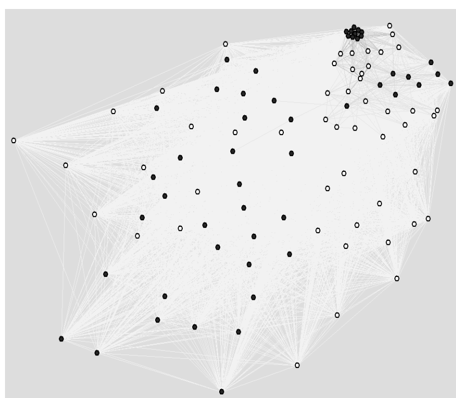

(f) 100 Experiments Results mixed population.

Fig. 1. Collaboration Graphs for different populations settings

population of just one experiment, and figure 1(d) represents the aggregated results amongst the set of 100 experiments with different tasks to fulfil.

In the first figure we can appreciate how collaboration is not concentrated in clusters, instead it is spread in many different combinations of agents. This suggests the correctness of the analysis performed in section 4.2, where it was stated that a conservative population does not necessarily converge into an stable state. This way agents move from coalition to coalition during the experiment, and the collaboration graph reflects that behavior by not showing any particular frequent cluster of agents, instead, the graph shows a cloud of sparse collaborations between the agents. Conservative agents optimise the ratio payoff/members keeping low the size of coalitions. The experiments reflect an average size of 2.3. The small size of coalitions, is a destabilisation fact that make them have similar scores. When coalitions have similar scores a single movement of an agent can change the whole payment scenario. However, some nodes in the center of the graph show slightly more frequent connections between them. This suggests that even though the system does not converge, there are successful agents that have some preferential attachments with other nodes with which they can create small coalitions with good scores, and so, it happens that in a dynamically changing environment they meet each other more frequently.

By observing figure 1(d) representing the aggregation of the data obtained across all the experiments, we can see that it keeps the same concentric structure, indicating that successful properties of central agents are kept. By the dark color in the edges, we can see that although conservative agents spread their collaboration, they usually cooperate with the same wide range of agents.

5.4 Mixed Strategies Dynamics Experiments

In these experiments, we mix 50 agents using the competitive strategy with 50 agents using conservative strategy. In order to create comparable results when playing two populations together, we create a symmetric population of 50 duplicated skills vectors. Figure 1(e) shows the *collaboration graphs* of the population in just one experiment. Figure 1(f) on the right hand side represents the results

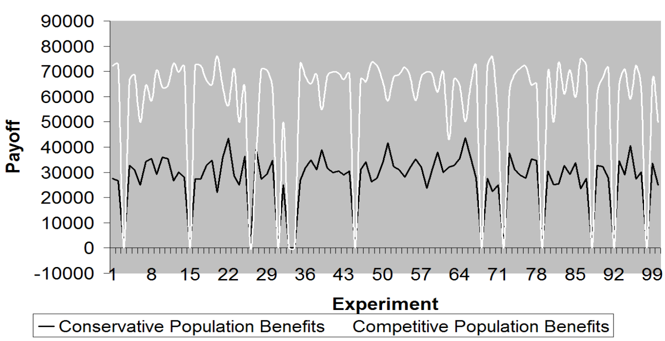

Fig. 2. Payoff gained by each population throughout 100 different tasks

amongst the set of 100 experiments with different tasks to fulfil. To differentiate members of each population in the graph, nodes representing conservative agents are in black color, and nodes representing competitive agents are in white color.

In the first graph we can see how a group with a majority of competitive agents establishes itself in the experiment, and show almost pure endogamic relationships. Apart from this, we observe that the rest of agents spread their collaboration with many different members of any population type. The average coalition size in this experiment is 5.9. This indicates that competitive agents boost the competitivity in the games, and large coalitions created by those agents become attractive also to conservative agents, as once the coalition has a certain competitive degree it becomes more profitable than small coalitions created by other conservative agents.

The figure showing results throughout the 100 experiments 1(f) reflects a very similar situation as in the one experiment graph 1(e). Both figures have a shape that reflects the mixture between clustering into frequent groups and spreading among many different ones. It is possible to appreciate a cluster of competitive agents, and very close to that cluster, reflecting frequent collaboration with the elements of the cluster, there is a set of conservative agents. Those agents do not stay in the cluster coalitions as frequently as competitive agents, because they are probably tempted by other coalitions offering less score and higher payoff, but they are probably good enough to be accepted in the coalitions when those coalitions become the most profitable option for a conservative agent at a certain time.

In terms of profit, as we can see in figure 2, the competitive population clearly outperforms conservative population. This is explained by the fact that competitive agents are the founder creators of the most competitive coalitions and remain on it as they cannot find any incentive to leave it. On the other hand, conservative agents that get attracted by a successful high score coalition might join it if its payout/members ratio is good enough, but they might eventually leave it seeking apparent opportunities in the market of smaller coalitions with higher benefits. This often turns out to be the wrong decision as smaller coalitions are more dynamic and subject to change (see section 4.2 for an explanation of this fact).

6 Related Work

There is other relevant work in literature that tackles heterogeneity of populations as a key point of the research. The most representative is [5] which studies how diversity within the agent population impacts on the quality of the coalitions that emerge. Conforth et. al. create a dynamic organization framework in which heterogeneity turns out to be a crucial element for problem solving tasks. In this work, heterogeneity is represented upon different initial values of agent's parameters such as their connectivity (interaction links), trading strategies and initialization states. In the present model, all agents know each other but are distinguished by different values in a fixed set of skills. This fact characterizes every individual by intrinsic properties that can be complementary, and lead to the system to have different dynamics from those showed in the cited work. Apart from using heterogeneity in the intrinsic properties of the agent, the presented model contemplates different strategic behavior for agents.

Another interesting example on the use of heterogeneity in coalition formation processes is the modelling of heterogeneous preferences. Some examples of this are [2] which models different preferences between leisure and work, and [4] which provides a method that considers the possibility of different evaluation functions to coalition structures.

From the point of view of the protocol used (RFP), the work presented is related to [11]. However, Kraus et. al. have radical differences in its use. Firstly, in their model, a number of tasks are issued and agents propose coalitions (from scratch) that are accepted or rejected by the affected members. In order to motivate agents to form a coalition, a discount factor is used. In the model presented here, there is just one Task at a time for which all the population compete for, but not only one coalition gets priced. Agents do not propose entire structures, instead they construct them by individual movements. In order to motivate agents to form good coalitions, the motivation instrument is the use of a decreasing payment function.

The presented model shares many characteristics with economic models, such as [2]. Some important characteristics are: the iterative nature of the processes, the episodical evaluation of the structures, and similar type of rational behavior in the agents, however there is an important difference that is the evaluation function. The evaluation function applied by Axtell et. al. (Cobb-Douglass) rewards a coalition (or a *firm*) independently of the rest of existing coalitions. In our case the reward of a coalition is partly dependant on its score, and on the score of the rest of the existent coalitions.

Traditionally Coalition Formation problems have been tackled as a "one off" event. In the present work we seek to go beyond this to consider what may happen in environments over time. Other important work in the same line includes [9,10,1].

7 Conclusions

From the analytical and experimental work presented, the following conclusions, applicable to the Iterative RFP domain, are drawn:

- Competitive strategies (score maximizing) outperforms conservative strategies (payoff maximizing) when symetric populations are played against one another. This provides a hint as to why competitivity emerges in certain societies to dominate other conservative behaviours.
- A pure population of agents running the competitive strategy for a give task, converges to a Nash equilibrium state in which no agent has motivation to move elsewhere given the coalitional structure created.
- A pure population of agents running the competitive strategy tend to create coalitions of size equal to the number of requested skills.
- A pure population of agents running the conservative strategy tend to create small size coalitions.

– In a pure population of agents running the conservative strategy, there is a inversely proportional relationship between the size of coalitions and the degree of dynamism.

Lastly, the long term aim of this work is to analyse the clustering structure of agents within an RFP population - and investigate how this affects performance. To this end, Collaboration Graph representation seems to present a useful method to analyze properties of the Coalition Formation process in such long running scenarios.

8 Ongoing and Future Work

Ongoing and future work includes a deeper research on the conditions under which competitivity dominates conservative strategies.

Other planned work includes study of individual characteristics that make agents extraordinary, and extract the important patterns that are exploitable with certain strategies such as those shown in the current paper.

It will be also of significant value to study the stability of results from a formal game theoretic perspective, as well as the game theoretic properties of the strategies we study.

Finally, we are investigating the use of other large scale network analysis techniques to have a deeper understanding of our domain. Such techniques include *Clique Overlapping*, and *t-core distribution.*

References

1. T. Arnold and U. Schwalbe. Dynamic coalition formation and the core. *Journal of Economic Behavior & Organization*, 49:363–380, 2002.
2. R. Axtell. The emergence of firms in a population of agents: Local increasing returns, unstable nash equilibria, and power law size distributions. Working Paper 3, Center on Social and Economic Dynamics, Brookings Institution, 1999.
3. V. Batagelj and A. Mrvar. Pajek-program for large network analysis. http://vlado.fmf.uni-lj.si/pub/networks/pajek/.
4. P. Cailou, S. Aknine, and S. Pinson. A multi-agent method for forming and dynamic restructuring of pareto optimal coalitions. In *Proceedings of the 1st conference on Autonomous Agents and Multi-Agent Systems, AAMAS'02*, 2002. Bologna, Italy.
5. D. Cornforth, M. Kirley, and T. Bossomaier. Agent heterogeneity and coalition formation: Investigating market-based cooperative problem solving. In *Proceedings of the 3rd conference on Autonomous Agents and Multi-Agent Systems, AAMAS'04*, 2004. New York, USA.
6. A. Dragulescu and V. M. Yakovenko. Exponential and power-law probability distributions of wealth and income in the united kingdom and the united states. Computing in Economics and Finance 2002 125, Society for Computational Economics, July 2002.
7. D. D. Gatti and C. D. Guilmi. Financial fragility, industrial dynamics and business fluctuations in an agent based model. *paper presented ad the conference Wild@Ace 2003, Turin, Italy, October 3-4*, 2003.

8. D. D. Gatti, C. D. Guilmi, E. Gaffeo, G. Giuioni, M. Gallegati, and A. Palestrini. A new approach to business fluctuations: heterogeneous interacting agents, scaling laws and financial fragility, 2003.
9. M. Klusch and A. Gerber. Dynamic coalition formation among rational agents. *IEEE Intelligent Systems*, 17(3):42–47, 2002.
10. H. Konishi and D. Ray. Coalition formation as a dynamic process. *Journal of Economic Theory*, 110:1–41, 2003.
11. S. Kraus and O. S. ang G.Tasse. Coalition formation with uncertain heterogeneous information. In *Proceedings of the 2nd Conference on Autonomous Agents and Multi-Agent Systems, AAMAS'03*, 2003. Melbourne, Australia.
12. C. Merida-Campos and S. Willmott. Modelling coalition formation over time for iterative coalition games. In *Proceedings of the 3rd conference on Autonomous Agents and Multi-Agent Systems, AAMAS'04*, 2004. New York, USA.
13. V. L. Smith. An experimental study of competitive market behavior. *The Journal of Political Economy*, 70:111–137, 1962.