

The effect of heterogeneity on coalition formation in iterated request for proposal scenarios

Carlos Mérida-Campos Steven Willmott

Knowledge Engineering and Machine Learning Group (KEMLG)
Universitat Politècnica de Catalunya (UPC)
Barcelona 08034, Spain

Abstract

This paper explores a general model of economic exchange between heterogeneous agents representing firms, traders, or other socioeconomic entities, that self-organise into coalitions to face specific tasks. In particular, the work addresses *coalition formation* problems in which many tasks are addressed to the same population over time in an iterative fashion – giving the agents the possibility to organise themselves for specific tasks.

In particular, the purpose of the paper is to describe the impact of task and skill heterogeneity in the population on the performance of the agents. Experiments are carried out for two common strategies from the economic world, namely *competitive* and *conservative* strategies. Results obtained show that competitive population outperforms conservative population, and that the heterogeneity degree has a direct effect on the advantage of the first strategy versus the second. We analyse experimental results by using a novel data mining technique called *collaboration graphs*.

1 Introduction

As defined by Whitesides and Ismagilov, a complex system is one whose evolution is very sensitive to initial conditions or to small perturbations, one in which the number of independent interacting components are large, or one in which there are multiple pathways by which the system can evolve. This paper shows results on the research of a complex system studied from an economic point of view. The specific complex system investigated is a general model of economic exchange between heterogeneous entities that form coalitions. These coalitions compete amongst themselves to solve a certain task in order to score as highly possible for the task. A third party thus evaluates the coalitions and ranks them according to their score in terms of skills for the task. Analysis of the macroscopic behaviour of such systems quickly makes it clear that, depending on the properties of the tasks and the agents, different long term Compatibility patterns are established amongst some groups of agents which often work together in successful coalitions. Intuitively also, in a population with a mix of skills, certain combinations of agents intrinsically work well together (i.e. they have compatible skill sets for certain types of task). Concretely, the paper shows that there exists a relationship between heterogeneity in the task definition and similarity of performance between competitive and conservative strategy. The more heterogeneous a task is, the more similar the outcomes of the two strategies. We also show that heterogeneity of the agents does not have the same effect on the outcomes of the strategies. We finally show that in a higher or lower degree, competitive strategy has a better performance than conservative one.

2 Background and motivation

The organizational paradigm of our economic model is Coalition Formation. These models have been extensively studied from a game theoretic perspective initially ([18, 10, 9] amongst others) and in the last years, from a multi agent systems perspective ([2, 3, 14] amongst others). Game theoretic

research of coalitional models focus basically on studying stability properties by analysing different stability solution concepts such as the core, or the kernel [10]. Dynamic Coalition formation further focuses on the study of coalition formation problem of complex systems, this area includes research from a formal mathematical perspective that model with markov chains coalition formation systems to study their properties with several strongly-couples degrees of freedom ([13, 1]). Multi agent systems research is also moving toward this direction in the research of coalition formation systems, and includes papers on the analysis of classical stability concepts of systems not constrained with the necessary rigidity of mathematical models ([12, 17]). The area includes research on a variety of different fronts, including problems such as finding stability properties between coalition structures [8] or finding solutions to a problem in a cooperative way [6]. In general, research on this area does not attach much importance to the topic of heterogeneity, and assume homogeneous individuals competing with homogeneous strategic profiles. An exception to this case is the work by [5], where the author highlights the importance of modelling heterogeneity. Concretely the work focuses on the role of externalities across actors in determining the population-wide behaviour. These externalities, as in our case, are in turn the source of interactions. The cited paper, however, does not focus on a concrete model and study it in deep as is our case, they study the general properties, applications and issues of what they call *Interaction Models*, from which our model is identified. Other papers such [14] or [3] use some degree of heterogeneity however focus rather on modelling different task valuations for each agent, agent capabilities and greediness levels in the agent’s expectations – rather than focusing on heterogeneity per-se.

Our model is driven by complementarities between agents and spillovers between coalitions, thus allows for the analysis of market imperfections in heterogeneous agent contexts.

In the market mechanism suggested in the present work, 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 [15] 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 that direction analyzing the dynamics of two concrete rational behaviours, a so called “Competitive Strategy” (agents try to be in a coalition with maximal competence) and a so called “Conservative Strategy” (agents try to be in a coalition with optimal competence / size ratio in order to maximize their benefits) within such an environment. These strategies are common examples of rational behaviour in economic markets. There are, indeed more complex strategies, but, as this paper argues these two strategies already prove sufficient to create two differentiated types of dynamics. In [16] the same two strategies are investigated in an isolated way and together in a general setup. This work extends the cited paper by focusing on the influence of the heterogeneity degree of agents in the performance results as well as on the collaboration structure they show up.

The paper is structured as follows: Section 3 explains the concrete market mechanism in use, the Iterative RFP Coalition Formation method. Section 4 explains the agent based system designed to model the iterative RFP coalition formation method, including the valuation function definition, and the agents’ strategies. Section 5 explains the concrete setup used for the experiments as well as the *collaboration graph* method used in the analysis of results. Section 6 analyses the results of the experiments with respect to the effect of heterogeneity on the strategies outcomes. Section 7 analyses the results on the effect of heterogeneity in the collaboration structures between agents. Finally, Section 8 wraps up the results obtained showing the conclusion of the present work.

3 Iterated RFP Systems

Request For Proposal Systems (RFP Systems from now on) are defined as those in which an individual or an entity submits an invitation (RFP) for providers of a product or service to bid on the right to supply that product or service to the individual or entity that issued the RFP. This type of systems is widespread in many different economic activities such as European Calls for project proposals in different frameworks, public contests for concessions of licenses to construct buildings, or to handle a particular marketing campaign and so forth.

Calls in real systems are announced and issued once or more times a year. They can be con-

sidered as iterative processes that repeat periodically over time. This way certain individual that participated in a certain call will likely participate in the subsequent one facing a related problem over and over, and adapting to the variations in the requirements and in the environment in order to be more competitive. Summarizing, the characteristics of the RFP System under study are:

- 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.
- Groups might change depending on the market situation.
- Agents are individual utility maximizers and they share the same preference, that is payoff maximization, but might have different strategies for maximizing their utility.

4 A model of a RFP system

A population $I = \{1, 2, \dots, n\}$ consists of a finite number of n individuals or *agents*. Agents compete amongst them given the requirements specified by a certain *task* T . A *partition* $\sigma = \{\sigma_1, \sigma_2, \dots, \sigma_p\}$ of the population I is a specification of p *coalitions* of agents $\sigma_i = \{\sigma_{i1}, \sigma_{i2}, \dots, \sigma_{im}\}$, where σ_{ij} represents an agent from population I forming part of coalition σ_i . By forming coalitions, agents are able to increase their competitiveness towards the specified task.

Agents have heterogeneous capabilities, thus having different performance levels in different skills. A finite number of k *skills*, indexed from 1 to k is set for which each agent σ_{ij} has a fixed value: $\sigma_{ij} = \langle \sigma_{ij}^1, \sigma_{ij}^2, \dots, \sigma_{ij}^k \rangle$. In this way it is possible to define a continuum of possibilities between agents that are *specialised* in the performance of a certain skill being unskilled for the rest of them, and agents that are *versatile*, being averagely apt for the performance of all the skills defined.

A Task T is specified by a set of k skill requirements. Each of the k skills have a degree of requirement. These requirements are modelled in the form of a number $T = \langle T^1, T^2, \dots, T^k \rangle$.

In a coalition, skills of agents are aggregated in such a way that each agent gives the best of itself in a join effort to create a group as competitive as possible under the requirements of the Task. The coalition has a value in each skill representing the aggregated effort of its members. The aggregation for every skill $j : 1 \leq j \leq k$ in the coalition is modelled in the following way:

$$\sigma_i^j = \max(\sigma_{ij}^l) : 1 \leq j \leq m \quad (1)$$

In this way, there is always one agent σ_{ij} in coalition σ_i that is providing the maximum value for every skill l . This is noted as σ_i^l . Each skill is considered as a necessary subtask for performing task T . 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 the one that performs it, specializing just in that subtask while others in the coalition perform the other subtasks.

Amongst the different possibilities for aggregating agent skills in a coalition, we have chosen the Maximization function, as we consider it 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.

4.0.1 Valuation Function

The value function determines the payout given to each priced coalition. The general properties of function desired are:

1. The function must be monotonically increasing with the size of the coalition. In this way, the function will measure how apt the coalition is to perform a task without considering how optimal is in terms of size. That will be considered in the payoff mechanism.
2. The function must be upper bounded. Corresponding to the idea of having a score for a perfect proposal and all suboptimal ones will have a lower score than that.
3. The function must be deterministic. Hence we model the fact that the evaluation criteria is common knowledge and exclude the need for objective valuation.

With these properties in mind one can formally define a score function $scr(\sigma_x, T)$ to measure how well agents in coalition σ_x perform together for accomplishing a task specification T without having into account the rest of coalitions.

As explained in section 4, Tasks are specified by a set of skill requirements, and each of the k skills has a degree of requirement. For the purposes of this scoring function, those requirement values will be treated as weights, modelling the fact that some skills are more important than other for the accomplishment of a certain task. The score of a coalition can be computed as the weighted sum of the best skill values for each skill that the coalition has. Formally:

$$scr(\sigma_i, T) = \sum_{l=0}^k (\sigma_i^l * T^l) \quad (2)$$

It is straightforward to check that this function complies with the three properties named above. While there are other possible value function the one chosen here reflects models used for coalition formation in a large part of the current literature.

4.0.2 Coalition's Reward

Once we have the endogenous value of a coalition, it is necessary to assign the payoff depending on its ranking. That constitutes the second step required for computing the general coalition utility (the coalition's payoff).

The reward function defined here, allocates wealth in a decreasing fashion depending on the score of the coalitions. To this end, a *ranking* function creates a strict order between coalitions by using $scr(\sigma_x, T)$ function as an ordering criterion, and ordering randomly those consecutive coalitions with equal score. This function is expressed formally as: $rank(\sigma_x, \sigma)$, and it returns the value of the strict order of coalition $\sigma_x \in \sigma$.

Coalitions are priced with an exponentially decreasing amount from the initial one following the order established by the rank function, and there will always be a fixed amount P of payoff to spread amongst the coalitions, independently on the number of coalitions existent. Formally:

$$CR(\sigma_x, \sigma) = \begin{cases} P/2^{rank(\sigma_x, \sigma)-2}, & \text{for the last competent coalition,} \\ P/2^{rank(\sigma_x, \sigma)-1}, & \text{for all the other competent coalitions.} \end{cases} \quad (3)$$

Amongst the different possibilities for assigning the payoff, 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 P . This is an important property for realistic CFP environments and also ensures one can compare the economic behaviour of different populations. The second reason is because this represents a power law distribution of wealth for which there is empirical evidence from Real Economies (see [7]).

As a simplification however, Agents within a coalition spread coalition's payoff evenly (and hence methods for differential allocation of payoff to individual agents are not applied here). While this removes some complex coalition formation issues it still forces a tradeoff for agents between coalition size (in order to be more competent, thus gain a higher score and being more group profitable), and coalition optimality (in order to share the payoff by a reduced set of members, and so being individually profitable).

4.1 Agent Actions

This sub-section defines the actions agents can take between different issuances of tasks. Each player's strategic variables are his coalition choice σ_j and a set of agents in this coalition $\phi_{jk} \subset \sigma_j$ to eliminate and optimise the coalition. This possibility of optimisation responds to the change of value that certain agents can experiment when in the coalition they are out-skilled by a new member. It is necessary to model this possibility as this model deals with heterogeneous individual capabilities (not only different strategies).

Time is discrete. In each period, a random subset of agents is chosen, and they are asked sequentially to take an action regarding its membership of a coalition, and the way this coalition should be optimised.

The new membership together with the optimisation proposal is not accepted straightforward, it is evaluated by the members of the target coalition. Just those actions accepted by a majority (more than the half) of members in the affected coalition, are performed. An agent that is requested to take an action can submit a finite number of requests in an specific order, in such a way that if an action is not accepted, next action is checked and so on. If none of its action proposals is accepted, the agent stays in the same coalition where it was.

Summarizing, the choices than an agent has are:

- Stay in the current coalition.
- Stay in the current coalition optimizing it by firing (expelling) one or more members.
- Leave the current coalition in order to join a different one.
- Leave the current coalition in order to replace one or more agents in a different one.
- Create a new coalition with itself as the only member.

Each of these actions potentially changes the coalition structures in the society and with many agents taking such actions each turn leads to an evolution of state.

4.2 Agent Strategies

This paper explores the effect of two possible agents strategies:

- *competitive strategy* (score maximizing)
- *conservative strategy* (payoff maximizing).

For the case of conservative strategy, the agent expects to obtain 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.

In both cases, an agent σ_{zi} constructs an ordered set of actions proposals if requested: $\psi_i = \langle (\sigma_a, \phi_{ab}), \dots, (\sigma_p, \phi_{pq}) \rangle$. Those actions are evaluated sequentially until one is accepted or none of them is. For the case of score maximizer agents, ψ_i contains all the possible proposals (σ_x, ϕ_{xy}) for which one of the following conditions is satisfied:

$$scr(\sigma_x - \phi_{xy} + \sigma_{zi}, T) > scr(\sigma_z, T) \quad (4)$$

or

$$scr(\sigma_x - \phi_{xy} + \sigma_{zi}, T) = scr(\sigma_z, T) \wedge |\{\sigma_x - \phi_{xy} + \sigma_{zi}\}| < |\sigma_z| \quad (5)$$

In words, a competitive agent's action proposal set contain every proposal that either improves the score of the coalition the agent is in, or keeps the same score while reducing the size. Analogously, for the case of conservative strategy, all proposals in ψ_i fulfil the same conditions but changing $scr(\sigma_x, T)$ by $CR(\sigma_x, \sigma')$, where σ' stands for the partition resulting from any modification has been done to σ in the first argument of the function.

The order between two action proposals (σ_a, ϕ_{ab}) and (σ_c, ϕ_{cd}) in the proposals set ψ_i for a given agent σ_{zi} is defined in such a way that $(\sigma_a, \phi_{ab}) > (\sigma_c, \phi_{cd})$ if and only if one of the following conditions are satisfied for the case of competitive strategy:

$$(src(\{\sigma_a - \phi_{ab} + \sigma_{zi}\}, T) > (src(\{\sigma_c - \phi_{cd} + \sigma_{zi}\}, T) \quad (6)$$

or

$$(src(\{\sigma_a - \phi_{ab} + \sigma_{zi}\}, T) = (src(\{\sigma_c - \phi_{cd} + \sigma_{zi}\}, T) \wedge (|\{\sigma_a - \phi_{ab}\}| < |\{\sigma_c - \phi_{cd}\}|) \quad (7)$$

Some proposals cannot be indexed with strict order, that is the case when having the score as a first ordering criterion and coalition size as a secondary ordering criterion, both proposals have the same characteristics:

$$(src(\{\sigma_a - \phi_{ab} + \sigma_{zi}\}, T) = (src(\{\sigma_c - \phi_{cd} + \sigma_{zi}\}, T) \wedge |\{\sigma_a - \phi_{ab}\}| = |\{\sigma_c - \phi_{cd}\}| \wedge ((a \neq c) \vee (b \neq d)) \quad (8)$$

In that case, a random order is established between consecutive proposals with equal value. Analogously, for the case of conservative strategy, the criteria used is the same but, again, using $CR(\sigma_x, \sigma_l)$ instead of $scr(\sigma_x, T)$.

Agents are not completely farsighted and neither they have information on the skills of other agents. The only information they have access to in addition to the their own skills and specification of the current task is:

- The potential score of their current coalition for the task.
- The precise outcome of any action they want to consider prior to its submission in terms of coalition score (however they have no information on the evolution of coalitions after its decision, nor the individual skills of agents in other coalitions).

The model considers a certain degree of randomness in the way agents are chosen to make a move. In every game, just a certain fixed number of actions are issued, and agents cannot make considerations on which agents will be elected to make a choice, nor what will be the order of the choices. Given these restrictions, we can consider agents in this model as myopically rational.

5 Experimental Set-Up

While the model presented in the previous section is quite specific (using a specific scoring function, pricing policy, set of strategies, etc.) it represents a good prototype for a variety of *iterative coalition formation environments* since it contains the major elements one would expect to see in such worlds: regular task arrival, agents slowly changing coalition structures over time to work with others and tension between coalition score and size due to payoff.

For all the experiments performed, a significant range of variations was tried and results found to be robust across them. These parameters for the results shown here are:

- Population size ($|I|$): 100
- Number of Skills (k): 10
- Total value of an agent ($\sum_{j=1}^k \sigma_{xy}^j$): 200 (it means, that the sum of every agent's skill will be 200 for all the agents)
- Total value of a Task ($\sum_{j=1}^k T^j$): 100 (it means, that the sum of every task's skill requirement will be 100 for all the tasks under test)
- Number of Games per experiment: 1000
- Number of Agent's decisions per game: 10 (it means, that the number of agents that will be asked to make a decision in one game is 10)

The rest of the parameters are variables related to the issue of heterogeneity discussed in this paper. Those parameters are:

- Heterogeneity of Population: value distribution amongst agents' skills
- Heterogeneity of Tasks: value distribution amongst task's required skills

All agents have the same probability of being selected to make a choice, that is $1/|I|$. The strategies of agents have been implemented by means of an exhaustive exploratory algorithm. When a competitive agent is requested to make a move, the algorithm checks all the possible actions of an agent. It checks the value of all possible join movements given the current situation, and for each option, checks every possible optimizations (see section 4.2) checking the score and the size of the resulting coalition. The algorithm constructs an ordered list of options that improve the situation the agent is in. Analogously, for conservative agents, it checks the payoff the agent would get instead of the score of the coalition.

In order to create experiments that represent a valid range of heterogeneous properties, three different tasks with different variability level (Sharpness) in their skills values distributions have been characterised. Namely the tasks types studied are: Low Sharp Task (*LST*), Medium Sharp Task (*MST*), and High Sharp Task (*HST*). *LST* has its 10 skills required with a similar skills value (mean value is 10, and stdev. is 5.). *MST* has not that similar skills requirements, and two of its skills are not required (mean value is 10, and stdev. is 10, two skills with value 0). *HST* has widely varying skills requirements, demanding just 4 out of its 10 skills (mean value is 10, and stdev. is 15). Experiments also include three different types of populations with different expertise level. Namely the categories studied are: Low Sharp Agents (*LSA*), Medium Sharp Agents (*MSA*) and High Sharp Agents (*HSA*). Each of the three populations has 100 agents, but populations are created in such a way that there are 50 different skill setups for agents using competitive strategy and the same 50 skill setups duplicated for agents using conservative strategy. *LSA* is a set of generic abilities, characterised by a low standard deviation in their skill values (mean value is 20 and stdev. is 10). The *MSA* set has moderate standard deviation in their skill values (mean value is 20 and stdev. is 20 for all of them). *HSA* is a set of specialised agents, characterised by their high standard deviation in their skill values (mean value is 20 and stdev. is 30 for all of them).

In order to check the effect of heterogeneity in task requirements and agents population, in this configuration a total of 270 experiments have been run. Sets of 30 different *LST*, 30 different *MST* and 30 different *HST* tasks have been generated to cover the space of tasks. The experiments for each set were then run for each of the different populations of agents (*LSA*, *MSA* and *HSA*).

5.1 Collaboration Graphs

An important aspect of the results shown later is the analysis of compatibilities of agents using collaboration graphs. In order to observe how the population complements, results are shown for a given set of tasks and agents using the following analysis:

1. Agents participate in the coalition formation process sequentially for the defined tasks, creating partitions of the population for as many games as defined.
2. After each game, the one-to-one relationships established between members of the population are saved in that game.
3. After the final game, the frequency of these one-to-one relationships are compiled across all games and used to generate weights in a population graph.

The data generated in this way contains information on which agents often work together for the task set (which is distributed across the space of possible tasks). As a result this can be used to analyse, how agents cluster together for the concrete parameterisation of the experiment. The mechanism used to determine clustering is an implementation of *Kamada-Kawai* [11] algorithm in *Pajek* [4]. This algorithm places nodes in a bi-dimensional space in a close position when they are

connected with links of relative high value. Graphs are created from data obtained by the clustering algorithm, and for the sake of the clarity, these also distinguish the frequency of cooperation between agents, not only by virtue of their position in the graph, but also with the colour of the arrows in such a way that they look darker when the relationships are more frequent, and lighter otherwise. Throughout the rest of the paper, these graphs are referred to as *Collaboration Graphs*.

6 Heterogeneity effects on strategic advantage

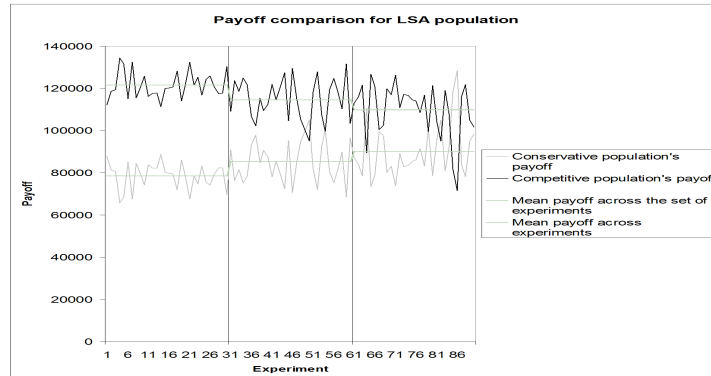
	LST	MST	HST
LSA	avg:707 stdev:184	avg:519 stdev:172	avg:195 stdev:155
MSA	avg:567 stdev:215	avg:455 stdev:230	avg:172 stdev:124
HSA	avg:581 stdev:155	avg:471 stdev:172	avg:212 stdev:163

Table 1: Average differences and standard deviation between the payoff obtained by each population of agents (competitive and conservative) for the set of 30 experiments using different tasks and agents with 3 different heterogeneity degrees. The mean value of payoff for every agent skill setup across the different tasks setups indicates the inverse relationship between task sharpness level and advantage of competitive versus conservative strategy


Competitive population (score maximization strategy) proves to be more effective throughout the whole series of experiments, outperforming conservative strategy (payoff maximization) on average. This fact is reflected in the total payoff obtained by each population in each experiment played. Figure 1 shows how the competitive population gets better payoff than conservative population in the majority of experiments. Table 1 summarizes these results. Values in the table are the average difference between what the competitive population obtained in an experiment minus what the conservative population obtained in the same experiment. The same table also reflects the standard deviation on these different values. The table shows numerically that competitive population outperforms conservative population for every type of heterogeneity combination.

Table 1 (as well as in Figure 1), shows that competitive populations do not behave homogeneously well in all the combinations. One can observe that there is a clear relationship between task heterogeneity and population profit. The more heterogeneous the task, the less advantage competitive population has. The key factor of this effect is coalition size. Coalitions formed when the task is heterogeneous have smaller size than when the task is homogeneous, because in the first case, some of the skills are not required, thus less expert agents are required to create a competent coalition. When the task requires many skills, competitive agents are willing to form coalitions which are as competent as possible, growing in size disregarding their benefits. The long term effect of this behaviour is that coalitions created by competitive agents turn out to be positioned in the top of the ranking, providing the highest benefits to their members. Conservative agents, in turn, optimise the competency/size ratio of the coalitions they form, thus reducing the size of the coalition as much as possible while keeping the same position in the competitiveness ranking. When the difference from the score of conservative agents to the subsequent coalition in the ranking is sufficiently high, conservative agent's coalition will disintegrate their structures, losing competence, and giving chances to other coalitions to outperform it to occupy its rank. This situation usually happens when coalitions are big in size. The long term effect of this behaviour is that conservative coalitions are unstable and outperformed by competitive coalitions in these scenarios.

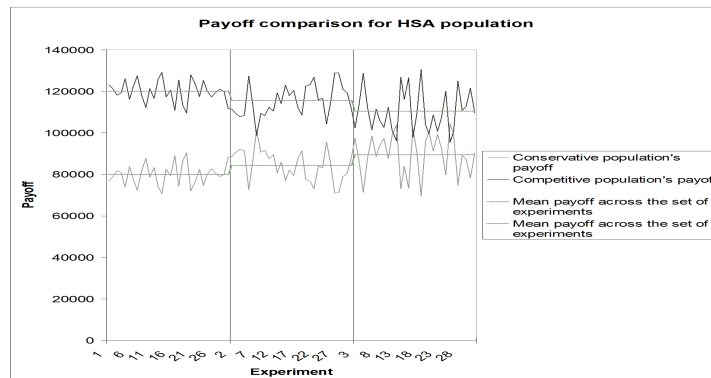
This disadvantage is diminished when tasks sharpness degree is higher, as coalitions need not to be big, hence relative distances between scores in the competency ranking of coalitions is lower. This will not allow conservative agents to optimise the size of the coalition they are in, behaving like competitive agents, and obtaining very similar profits.



(a) Payoff obtained for each strategy population in a LSA skill setup for every experiment in each one of the three tasks setups (LST, MST and HST).




(b) Payoff obtained for each strategy population in a MSA skill setup for every experiment in each one of the three tasks setups (LST, MST and HST).




(c) Payoff obtained for each strategy population in a HSA skill setup for every experiment in each one of the three tasks setups (LST, MST and HST).


Figure 1: Payoff distribution for experiments using each one of the three agent skills setup. The mean value of payoff for every task setups indicates the inverse relationship between task sharpness and advantage of competitive strategy versus conservative.




(a) Collaboration graph corresponding to the setup LST-LSA




(b) Collaboration graph corresponding to the setup LST-HSA




(c) Collaboration graph corresponding to the setup MST-LSA



(d) Collaboration graph corresponding to the setup MST-HSA



(e) Collaboration graph corresponding to the setup HST-LSA



(f) Collaboration graph corresponding to the setup HST-HSA

Figure 2: Collaboration graphs of 30 experiments each, of a certain agent and task sharpness. Black vertices represent competitive agents and white nodes represent conservative agents. Comparing graphs of differently Sharp tasks we see a direct relationship between task sharpness and pattern of collaboration frequency

7 Heterogeneity effect on collaboration structures

The system under study generates different dynamics in the agents that make them create more or less similar coalitions throughout the whole set of experiments. The experiments revealed that it exists less stability and less similarity in coalitions created in the low Sharp tasks scenarios compared to those created in the high Sharp task ones. This could be counter intuitive given the fact that lower Sharp tasks are more similar to each other than higher Sharp tasks. This phenomenon, as in the case of the relationship between advantage of competitive population and sharpness level, is independent of the type of agent sharpness used.

Figure 2 shows the collaboration graphs for two different agents populations across the three different tasks sharpness. One can see that the higher is task sharpness, the more frequent are the relationships between agents (darker arrows between nodes). For HST graph case, it is also possible to distinguish different clusters of frequent collaboration between agents. The reason, as in the previous case, is the reduced size of coalitions formed in the HST scenarios. In the low sharpness setup a clear cluster of competitive agents appears. These agents are the best equipped agents of the population, and they form a stable and competitive coalition that finds no competence in the conservative population because of the disintegration effect explained in previous section. For lower ranking coalitions, competitive agents find more difficulties in stabilising themselves because there are competent conservative agents, that instead of forming a competent coalition, de-stabilise the rest of coalitions optimising their size and re-combining with different groups. For higher task sharpness levels, given that the disintegration effect in conservative agents is cancelled out by the reduced size of the coalitions, they create more stable collaboration structures.

8 Conclusions

The work presented is based on a model of coalition formation using a RFP protocol for the type of iterative coalition formation environment often found in competitive real world scenarios. While simple, this model can be used to study agent's behaviour in such systems that can be represented by the model abstraction. The results specifically address the effects of heterogeneity using two different coalition change strategies. From the experiments and the analysis performed we draw the following conclusions for these environments:

- Competitive strategy is more advantageous on average than conservative strategy independently of the heterogeneity level in the agent skill setup and the task skill requirements.
- There exists a direct relationship between task heterogeneity and similarity of results obtained by competitive and conservative strategy. I.e. the more heterogeneous is the task definition, the more similar will be the outcomes of the two populations
- Heterogeneous tasks make conservative strategy behave similarly to competitive one given that the optimization opportunities are scarce.
- There is no relationship between agent skill setup heterogeneity level and advantage or disadvantage of none of the strategies.
- It exists a direct relationship between task heterogeneity and frequency of collaboration patterns.

Above and beyond this, the experiments also show the clear result that in such populations the intrinsic compatibilities between agent skills emerge as important structuring criteria even if the agents do not specifically know each others skill sets. While this seems trivially true, it is rarely exploited in coalition formation work to date. In future work, the hope is to develop a more complete model of the evolution of coalition structures in such populations over long running episodes of iterative or continuous varied task arrival.

References

- [1] Tone Arnold and Ulrich Schwalbe. Dynamic coalition formation and the core. *Journal of Economic Behavior & Organization*, 49:363–380, 2002.

- [2] Robert Axelrod and D. Scott Bennett. *The complexity Of Cooperation-Choosing Sides; A landscape Theory of Aggregation*, chapter 4. 1997.
- [3] Robert 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.
- [4] V Batagelj and A. Mrvar. Pajek-program for large network analysis. <http://vlado.fmf.uni-lj.si/pub/networks/pajek/>.
- [5] Lawrence E. Blume and Steven N. Durlauf. *The Interactions-Based Approach to Socioeconomic Behavior*. Social Dynamics. The MIT Press, 2001.
- [6] David Cornforth, Michael Kirley, and Terry 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.
- [7] 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.
- [8] Domenico Delli Gatti and Corrado Di 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*.
- [9] J. Greenberg. *Coalition Structures*, volume 2, chapter 37. 2004.
- [10] Sergiu Hart and Mordecai Kurz. Endogenous formation of coalitions. *Econometrica*, 51(4):1047–64, July 1983.
- [11] T. Kamada and S. Kawai. An algorithm for drawing general undirected graphs. *Inf. Process. Lett.*, 31(1):7–15, 1989.
- [12] M. Klusch and A. Gerber. Dynamic coalition formation among rational agents. *IEEE Intelligent Systems*, 17(3):42–47, 2002.
- [13] Hideo Konishi and Debraj Ray. Coalition formation as a dynamic process. *Journal of Economic Theory*, 110:1–41, 2003.
- [14] S. Kraus, O. Shehory, and 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.
- [15] Carlos Merida-Campos and Steven 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.
- [16] Carlos Merida-Campos and Steven Willmott. Agent compatibility and coalition formation: Investigating two interacting negotiation strategies. In *Proceedings of the AMEC Workshop*, 2006. Japan.
- [17] Tuomas W. Sandholm and Victor R. Lesser. Coalition formation among bounded rational agents. Cmpsci technical report 95-71, Computer Science Department, University of Massachusetts at Amherst, 1995.
- [18] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behaviour*. Princeton University Press, 1944.