

Modelling Coalition Formation over Time for Iterative Coalition Games

Carlos Mérida-Campos
Universitat Politècnica de Catalunya
Modul C6, Campus Nord, UPC
Barcelona, Spain, E-08034
dmerida@lsi.upc.es

Steven Willmott
Universitat Politècnica de Catalunya
Modul C5, Campus Nord, UPC
Barcelona, Spain, E-08034
steve@lsi.upc.es

Abstract

Coalition formation problems arise when groups of agents need to work together to achieve tasks in an environment – such as bidding for a contract or bulk buying goods. The work presented here shows how current theories for coalition formation can be combined with notions from iterative games to cover cases where populations of agents must solve coalition problems many times – modelling a long series of coalition games rather than just a single one.

The paper includes a problem formulation for iterative coalition games, experimental results for a simple coalition game world demonstrating how strong coalitions can emerge over time even from basic strategies and a discussion of the interactions between different strategies over time.

1 Introduction

Coalition formation problems occur whenever groups of agents must work together to solve one or more problems but must select the groups themselves from amongst a larger population of agents. Typical examples of such problems are forming consortia for contract tenders, consumers or suppliers teaming up to create economies of scale. In the general case such problems are computationally intractable (NP-hard) since each agent may have different skills and/or goals – leading to a large number of possible divisions of the entire population into subgroups – and further whichever process is used to form coalitions may in some instances be manipulated by agents to serve their own self interests (i.e. not leaving established coalitions to allow a more appropriate coalition to form).

While a wide range of different approaches to the coalition formation problem have been developed (diverse examples include using mathematical models [Lerman03] or heuristic

approaches [Kraus03]) one area that is less explored is what happens when the same population of agents (or one that is changing gradually over time) engages in many coalition formation episodes over time – either in sequence or in a continuously overlapping manner. That is new sets of tasks arise at regular intervals or in parallel and for each new task or set of tasks coalitions of agents are needed to address it. This type of scenario can be found readily in environments such as international commerce, bidding for government contracts or continuous auctions. In each case players (companies, contractors, bidders) may establish alliances for long period, cover multiple markets the state of the world changing over time between each coalition formation episode may affect decisions within the individual episodes.

In this paper we explore this type of environment in a simple manner by combining notions from the coalition formation problems with notions from iterative game theory which provides a rich literature for evaluating rational behaviour of individuals in repeated game problems.

In particular:

- We define the notions of iterative and continuous coalition games.
- Demonstrate how simple strategies applied by agents can lead to the emergence of strong coalitions in simple iterative coalition worlds.
- Show a number of interesting parallels between standard game theory results and strategic behaviour in iterative coalition worlds

Sections 2 and 3 respectively define the problem and experimental setup, Section 4 outlines major results, Section 5 provides discussion and Section 6 concludes the paper.

2 Problem Definition

Following descriptions such as those found in [Sandholm99] and [Klusch02] a coalition formation problem can be defined at its simplest as:

- Given a population P of agents and a list of tasks or goals T .
- Select subgroups of agents S_1, S_2, S_3, \dots of P to address each of the tasks in T .

A variety of problems then address the properties of the subgroups (stability, maximum social welfare, pareto efficiency, etc.), their behaviour (how any payoff is split or how the coalition is maintained to complete the task) and how the coalitions can come into being (the amount of information available, whether agents are cooperative, whether agents can be part of multiple coalitions and so forth).

In particular the *coalition formation problem* deals with the process of finding a set of combinations of agents which best solve given tasks in a given problem episode T (on set of tasks) for a given population of agents. From now on defined as a single *coalition game*.

Iterated coalition formation extends this notion by drawing the literature of *iterated games* used in game theory to lead to an *iterative sequence of coalition games*:

- Given a population P of agents and an iterated sequence of lists of tasks/goals T_1, T_2, T_3, \dots
- Let each task T_i correspond to a single coalition game G_i .
- Select for that G_i subgroups of agents $S_{i1}, S_{i2}, S_{i3}, \dots$ of P to address each of the tasks in T_i .

Intuitively this means that a population of agents persists over time to experience a series of coalition games – one after the other. (It should be noted that this is substantially different from existing literature on iterative approaches to coalition formation itself such as [Konishi03] – which are in fact techniques for solving a single coalition game – see Section 5.)

A further logical extension of iterative games is to *continuous coalition games* where individual games no longer happen in sequence but occur randomly in time and may overlap with one

another – forcing agents in the population to assess the tradeoffs between participating not only in particular coalitions but also between active concurrent games. The continuous case is also of interest, however the experiments in this paper cover only the iterative case.

3 Game Set-Up and Experimental Design

As noted, the definition given in the previous section covers a huge range of systems (or worlds). In particular: different levels of information may be available in different worlds, agent populations could be fixed between rounds or change, the type and number of tasks as well as the payoff distribution between coalitions or coalition members could all take a wide range of forms, finally the nature of the players can be cooperative or self-interested

Further, in terms of agent strategies – any coalition formation approach for single coalition game could be employed by agents in each round.

In this sense, the *iterated coalition formation problem* can be seen as a “macro-level” or “meta-level” problem at the level above problems in each *coalition game*.

To focus attention on these macro-level effects the experiments describe here adopt a very simple game world. In particular:

1. A static population of agents is assumed which may form coalitions of up to a fixed maximum size of members.
2. Problems arise in the form of contract tenders issued by a body such as a government – with the strongest coalition matching the profile selected each round for the particular contract.
3. Agents have no information about one another – knowing only their own skills and when the coalition they are part of wins in any given round.
4. Agents have only two actions they can carry out in a particular round: *stay* with the coalition they belong to or *leave* it.
5. Unassigned agents (e.g. those leaving coalitions) are randomly assigned to one of the other coalitions as long as the size maximum is not exceeded. With a certain probability they may also create a new coalition.

6. Payoff in winning coalitions is split evenly by participants in those coalitions, only one coalition wins in each episode.

While this makes for a very restricted game world it provides for at least a rough approximation of real contracting worlds: the number of companies in a particular sector is in generally relatively static in the medium term (approximating 1); while there may be information on the skills of others – it may well be inaccurate – especially if agents are involved in multiple bids and resources are spread (approximating 2); random assignment gives agents very little control and could be replaced by richer set of actions – however given that often entering a coalition involves complex negotiations and at any one point in time many coalitions may already have a full complement randomness provides at least a rough approximation of the process (approximating 5); lastly evenly splitting the payoff has been shown to be inequitable in general in other work [Sandholm99] however since agents have no information about the size of coalitions they may join or its members payoff division is dominated by the benefit of being in a winning coalition (approximating 6).

Given the significant simplifications each of these assumptions makes, each of these parameters could be extended to study the effect on the agent populations and overall outcomes.

4 Experiments

A range of experiments were carried out to test different hypotheses in the system. The same underlying configurations were used for the game world:

- A fixed set of 100 agents with abilities randomly distributed across 10 different skills. Each skill of an agent may be assigned a positive integer score or zero. Agents are each assigned a number N1 of skill points randomly distributed across their skills.
- A fixed repeated task (the same task used in all episodes/games) with requirements N2 skill points distributed randomly across the same 10 different skills.
- One winning coalition selected each game receiving a fixed payoff split evenly between team members.

- The winner selected using the following function defined in Box 1.

- A task **t** is **fulfilled** by a coalition **C** if the value for all the skills in **t**, is less than or equal to the maximum value among the members of **C** for the same skill.
- For a coalition **C** which fulfils **t** the **surplus** of **C** is the sum of all the skills in the task with value > 0 of the difference among the maximum skill value of an agent in the coalition and the required skill value of the task
- For a coalition which does not fulfil **t**, the **deficit** of **C** is defined as the subtraction of the difference among the maximum skill value in the coalition and the skill value for all the skills required in the task.
- The **winner** is chosen as: A) if there is a coalition that fulfils the task is the coalition with maximum surplus, or B) if there is no coalition that fulfils the task is the coalition with minimum deficit

Box 1: Winner determination function per game:
intuitively this picks the best coalition which meets all criteria OR if there is no such coalition the best coalition overall.

The game setup therefore uses a single unchanging task which is issued many times (once per game) and a winning coalition chosen. The winner determination method prefers the best coalition meeting the criteria for the task or that which has the least shortfall.

In addition, given their potential effect on the nature of the system, the following parameters were set separately: 1) the maximum coalition size (with tests at size 3 and 6), 2) the uniformity of the skill distribution in the population and task (ranging from regular distributions – most skills within 20% of the mean, to highly irregular – skills up to 80% from the mean), 3) The degree of difficulty that an average coalition of agents would have for fulfilling a task measured in terms of the ration between (N1*Max Coalition Size) and N2 – ranging from 2 and 4. The combinations of these three variables leads to 8 different game worlds ([3, irregular, difficult], [6, regular, ...] etc.).

In each of these worlds, combinations of a range of agent strategies were tested. In particular, agents could apply one of the following strategies:

- Random: each round randomly decided whether to leave their current coalition.
- Stay: always staying with their current coalition.
- Leave: always leaving their coalition
- Stay if Win (SIW): stay with their current coalition if it received a payoff in the last round.
- Stay in Win in One of the Last Two (SIW-2): stay with their current coalition if it received a payoff in either of the last two rounds. (Strictly if it did not loose with this coalition in either of the last two rounds to cover the case where the agent recently moved to the coalition.)
- Stay if All Stay (SAS): Stay in the current coalition only if all the other agents in that coalition stay.

Experiments were run with pure populations of 100% of each agent strategy and mixed populations of 50%:50% ratios of each combination of agents. This leads to 21 different experimental combinations – only some of these are reported here. (Note that in fact a number are redundant since the activities of the agents interact in such a way as to be equivalent to one of the pure populations.)

All results given in Section 4 are for worlds with a max coalition size of 6, irregular skill distributions and Difficult problem settings. 20 different skill distributions were generated for the 100 agents and one run of 500 games each population, carried out for each of these 20 different distributions. In each game the agents are informed if they have won or not, and then post their action for the next game.

In addition to the experiment runs for each population of skill distributions, game world parameters and agent strategies a brute force algorithm was also used to determine the best possible coalition given for the current experiment (i.e. the best coalition based on skills, maximum coalition size and task requirement). This value provides an upper bound on the total score of any coalition for a particular experiment and can be used to compute the Smith's Alpha measure [Smith62] for convergence during the simulation.

4.1 Pure Populations

The first set of experiments tested pure strategies amongst the agent populations – i.e. where all agents in a population play the same strategy

4.1.1 Alpha Convergence Experiments

Figure 1 shows the results of this metric for each one of the pure populations for the score achieved by the final coalition in each 500 round experimental run. The closer the value is to zero the closer the coalition winning in the last round was to the optimal coalition. Several things stand out from this:

1. Performance v's to optimal (zero) is heavily dependent on the skill distribution in the population (different in each numbered experiment on the x-axis).
2. The performance of each type of strategy seems equally dependent on the skill distribution in the environment.
3. Some strategies (notable SIW and SIW2) clearly converge routinely more effectively than others.

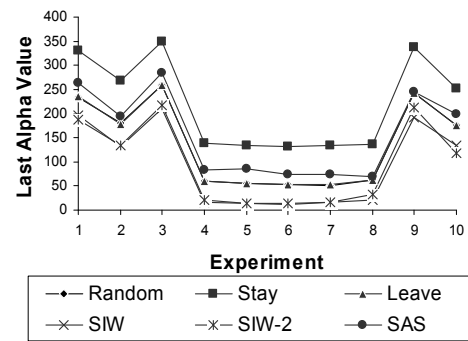


Figure 1.: Alpha Values for the pure populations in the first 10 experiments (results for the remaining 10 are similar).




Figure 2.: Amount of Possible Optimal Coalitions for each of the 10 first experiments set-up..

Upon analysis the dependence initial skill distribution is correlated to be due to the number of possible optimal coalitions in each randomly

generated game world. This is shown in Figure 2. In practice some game worlds create situations where many thousands of coalitions can produce the optimal result (or close to it) and some only very few. The differences in performance are therefore significant. However as shown in Figure 1 the relative performance of different approaches tends to hold.

These results also provide us with an initial idea of the performance of each strategy:

- Stay performs particularly poorly: this is due primarily to the fact that by itself it has no mechanism to adapt and hence over time improve on the initial set of single agent coalitions formed. A second factor is that coalitions do not grow to maximum size allowed by the game world – leaving many small, ineffective individuals.
- SAS performs better than stay but also poorly: as with stay SAS does not adapt – however in the first round agents by themselves group together to before settling into fixed coalitions (hence improving slightly on stay).
- Leave and Random Performs definitely poorly than SIW and SIW-2 and its unstable nature makes impossible to stabilise the possible optimal (or suboptimal configurations that can be found), but the fact that they are continuously exploring the space makes it better than Stay and SAS
- SIW and SIW-2 both outperform all the other strategies with little to choose between them (not statistically significant difference).

4.2 Mixed Populations

After the baseline experiments on pure population simulations were carried out on 50%:50% mixed strategies – that is assigning half of a population to play one strategy. 2 experiments were carried out in each case – with the strategy assignments switched around in each case.

4.2.1 Alpha Convergence Experiments

In terms of convergence to optimal values, in general results worsen – with most population combinations resulting in convergence towards the optimal coalition score being reduced. This trend was not followed in just three cases: mixing SIW with SIW-2, mixing SAS with SIW and mixing SAS with SIW-2.

In these three cases, the results of the alpha metric show very similar values to the best of the two strategies by itself. That is the stronger strategy interacted positively with the weaker one to speed up the convergence of both. An example of this is shown in Figure 3: SAS-SIW-2 performs just as well as SIW-2 does by itself. In this case it appears that SIW-2 agents:

- Move around gradually to improve overall winning coalitions
- Break up weaker coalitions after some time if they are not winning (by leaving and triggering the SAS agents to leave also).

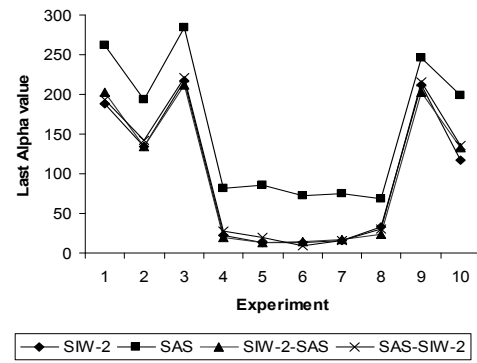


Figure 3.: Alpha Values for the pure SIW-2/SAS and SAS/SIW-2 populations in the first 10 experiments .

4.2.2 Payoff Comparisons

In addition to convergence measures it is interesting to compare the cumulative payoff received by each half of the population in each given game: i.e. to assess which strategy is dominated the other in terms of payoff gathered. Figure 3 shows in percentage terms the benefit of each strategy over each other summed over all rounds of all games




Figure 4: Income for each population mixed with each other population.

Comparing the incomes, the following results stand out:

- Random and Leave have few opportunities of getting paid (they are exploited) – particularly when mixed with SIW, and above all with SIW-2.
- Even though dominant on average, SIW and SIW-2 can still have worse results than Random and Leave in some experiments (See Figure 5 – where this occurs in population 12). This is due at least in part to the influence the concrete setup of skills. In this way, for some experiments, agents that form the community of Random could be much better for the prescribed task than those in the SIW or SIW-2 community, and even behaving not efficiently, the random coalitions outperform the best possible SIW and SIW-2 coalitions.

Although having worse results than Random and Leave as pure strategy, SAS gets better results than both of them in the mixed population experiments. This appears to be due to the fact that SAS agents are disrupted by Random and Leave agents (and hence do not become stuck). SAS perform better in general because they are more than Random/Leave, and so, it can happen that an stable coalition winning over time could contain members of SIW/SIW2 and SAS.




Figure 5: Payoff for each experiment having mixed Leave and SIW populations

4.3 The “Exploitation” of Leave by SIW and SIW2 Case

When choosing strategies SIW-2 was originally introduced as an approximate analogy of Axelrod’s Tit-for-Two-Tat’s (T42T) strategy [Axelrod81] and SIW as an analogue of Tit-for-Tat (T4T). However the results initially seem counterintuitive:

- Since SIW and SIW-2 perform nearly equally well in pure populations (where as T42T performs well in an environment of T2Ts).
- While both SIW and SIW-2 significantly dominated Random and Leave strategies (both gaining a large share) SIW-2 did so significantly more successfully.

The second result seems to be explained by the fact that SIW and SIW-2 are the only strategies that can be stable among the games of the experiments once they have formed a winning coalition, and so, if a population of SIW/SIW-2 has created a coalition with good score, they will be receiving the payoff as long as a better coalition does not appear. If a better coalition appears and this coalition is pure with members of SIW/SIW-2 then the advantage is enlarged, and the defence against new better coalitions is reinforced. On the other hand, if the coalition that improves the current one has some non stable member (from Leave or Random populations) those non-stable members will soon leave the coalition without making use of their advantage – thus disrupting the original strong coalition without establishing a new stable one.

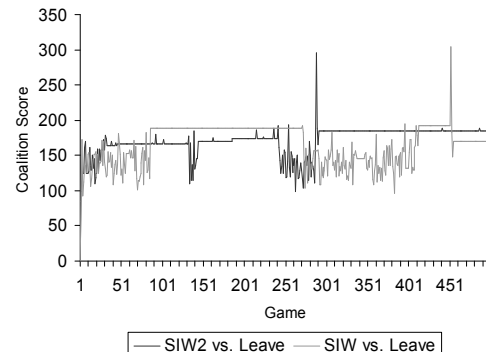


Figure 6: Score of the winning coalition during two experiments: using SIW2/Leave populations and SIW/Leave populations

The advantage SIW-2 has over SIW in this situation is clearly shown in Figure 6, which plots to score of the winning coalition in each round in a specific game. Each population is disrupted at least twice however SIW clearly takes longer to recover from disruptions – as when coalitions are weak they are much more easily disrupted by short term unstable coalitions.

4.4 Different World Configurations

As noted earlier in the section a range of parameters such as maximum coalition size affect the way the game world works. In general results shown here for the [6, irregular, difficult] case hold in other game configurations with the following observations:

- Using more regular task/skill as expected distributions smoothes out the variation which occurs due to different skill distributions.
- Using smaller coalition sizes convergence to high scoring coalitions is (as expected) considerably quicker than with size 6 coalitions.
- Using a smaller size of the coalitions permitted, the effect shown in Figure 5 is also more prominent. This appears to be because it is statistically faster to randomly a coalition that can break the hegemony of a winning SIW coalition (if it exists).

5 Discussion and Related Work

The systems described here are interesting from a number of points of view:

- The experiments show agents exploiting knowledge about past success between coalition games. In the absence of information and control in each individual coalition game certain strategies clearly exploit that limited knowledge they have on what happened in the last game to gain an advantage: this indicates interesting possible interplay between macro and micro level factors (information about players, information from history).
- Although there is no explicit notion of punishment in the world, the analogous of classic iterative game theory strategies clearly work well relative to simpler strategies: rewarding failure with disloyalty. As in the prisoners dilemma this extricates agents from poor situations but allows strong cooperation to materialise.
- Certain strategies interact positively and negatively (for example SAS and random or SWI-2 and SAS).

Most coalition formation work to date addresses the problem of forming a single set of a coalition to address a task / set of tasks for example and further often assumes that agents have significant information regarding one another's skills

(exceptions to this include [Kraus03]) or are interested in achieving group utility rather than maximising their own.

The work presented here provides a macro level framework to situate individual coalition formation approaches – opening the possibility of mixing strategies used within a game with those applied across several games. This likely to become particularly interesting if small amounts of information such as reputation indicators were introduced. (A simple way of doing this would be to post a “money list” declaring which agents had earned how much so far and allow agents to use this information to decide whether or not to stay / leave from coalitions.)

It is also important to note that the approach is different to iterative coalition formation approaches such as [Konishi03] which address coalition formation for a single episode. Similar elements are present between games as agents converge to a stable / dominant coalition over several games – however in this case each intermediate step is itself a game (and worth winning). There are also similarities to work on dynamic coalition formation [Klusch02, Soh02] which addresses how coalitions can be maintained over time in the face of change once they are formed.

From a game theory perspective related work uses a similar formulation to Axelrod's original iterated prisoners dilemma games [Axelrod81]. Providing agents a series of opportunities to work together – simple strategies such as Tit-for-Tat and Tit-for-2-Tats also appear to have analogies. Even though there is no direct reward/punishment structure – results such as the disruptions of stable coalitions in Section 4.3 that similar stability criteria apply – the forgiving nature of SIW-2 provides benefit over SIW.

6 Conclusions and Future Work

In conclusion, the iterative coalition games studied here represent a natural fusion of theory from iterative games and coalition formation problems – representing populations of agents working together overtime to in a limited form of persistent economy. The systems tested demonstrate:

- Complex dynamics even in simple strategies and clear gain in utility in agents exploiting

inter-game knowledge and strategies changing behaviour when exposed to different social environments

- A non-trivial relationship between game solitary player game strategies in iterative games such as the prisoners dilemma and a games which require larger scale collaboration.

The cases investigated remain very simple, however from these baseline set of results we see a large number of possible lines of investigation:

- Allowing multiple simultaneous tasks in each episode and allowing tasks to change over time (changing the optimal coalition and forcing agents to consider which tasks to bid for).
- Allowing tasks to arise continuously and span longer periods of time – forcing agents to consider the opportunity cost of participating.
- Allowing the population of agents to change between rounds (modelling a changing set of corporate players).
- Allowing more information to accrue in the environment such as past performance of particular agents or the scores of the winning coalitions (leading to new strategies and social behaviour).
- Allowing richer coalition formation mechanisms in each coalition game (applying existing techniques).
- Deeper investigation of the equivalences between iterative game theory results and iterative coalition game worlds.
- Allowing strategy and or skill change mutate in agents over time to would create an evolutionary coalition game world to investigate the dominance and stability of strategies with respect to one another.

Above all the formalization of iterated continuous coalition games appears to provide a useful perspective for bringing together work from a number of areas – iterative games, dynamic coalition formation and individual coalition formation techniques. Each of the extension listed is likely to change system dynamics in a significant way and bring in results from the related fields of Coalition formation (analysing the effect of external information on each coalition game), game

theory (bringing factors from individual coalition formation events into iterative and continuous contexts) and finally social simulation (understanding complex social structures which may emerge).

In terms of real world systems some of the above cases are also interesting – bearing a relationship to open economies in which companies compete for contracts and emerging markets over time: forming and dissolving alliances, balancing resources between markets and exploiting the varying levels of market/competition information that the environment allows.

7 References

- [Axelrod81] Axelrod R., Hamilton W. “The Evolution of Cooperation” Science, 211:1390-1396, 1981.
- [Klusch02] Klusch M., Gerber A. “Dynamic Coalition Formation among Rational Agents”.IEEE Intelligent Systems, 17(3):42-47, 2002
- [Konishi03] Konishi H., Ray D. “Coalition formation as a dynamic process” (2003). Journal of Economic Theory 110, 2003.
- [Kraus03], S. Kraus, O. Shehory, G. Tasse, “Coalition Formation with Uncertain Heterogenous Information”, Proceedings AAMAS’03, 2003.
- [Lerman03] Lerman K., Galstyan A., “Agent Memory and Adaptation in Multi-Agent Systems” (July 2003). Proc. AAMAS’03
- [Lerman00] Lerman K. Shehory O. “Coalition Formation for Large-Scale Electronic Markets” Proc. ICMAS 2000.
- [Sandholm99] Sandholm T., “Distributed Rational Decision Making” in Multiagent Systems: A Modern Introduction to Distributed Artificial Intelligence, MIT Press, Cambridge, Mass. 1999.
- [Smith62] Smith V. “An Experimental Study of Competitive Market Behaviour”. Journal of Political Economics, 70, 111-137, 1962.
- [Soh02] Soh L., Tsatsoulis C., “Real-Time Satisficing Multiagent Coalition Formation” Proc. AAAI Workshop Coalition Formation in Dynamic Multiagent Environments, AAAI Press, Menslo Park, Calif. 2002.