

Individual Versus Group Rationality: A Coevolutionary Approach to the Beer Game

Hongliang Liu, Enda Howley and Jim Duggan

Department of Information Technology, National University of Ireland, Galway

Email: liuhonglianger@gmail.com, enda.howley@nuigalway.ie, and jim.duggan@nuigalway.ie

Abstract

In this paper, we examine the evolution of strategies in the Beer Distribution Game (BDG). This is a well known game which has many parallels with supply chain optimisation problems. This paper explores the strategies used by individuals participating in this game. The issue of bounded rationality is addressed in detail throughout our simulations. This paper presents a new and alternative approach to evolving strategies for the BDG. This includes a co-evolutionary architecture which provides a means of evolving strategies across the various sectors of the BDG. Our results identify the best performing or fittest strategies for the BDG when two alternative fitness approaches are used. We show the significant implications for individual strategies in the BDG when agents are influenced by greater levels of bounded rationality. Our conclusions indicate the implications of individual and group rationality on the BDG.

1 Introduction

The BDG is a well-known model in the system dynamics literature (Forrester, 1961; Sterman, 1984). This production distribution game has been widely used to examine human decision making behaviour and inventory management optimisation (Sterman, 1987; Thomsen et al., 1992; Knolmayer et al., 2007). The traditional game normally involves four individuals, a retailer, a wholesaler, a distributor and a manufacturer. Each of these individuals faces a decision making challenge involving how they manage their current stock inventories. Each individual in the game seeks to minimise their total cost by managing their inventories in the face of uncertain demand. It has been shown that this simple game provides complex and often non-linear dynamics due to feedbacks and time delays. It has also been shown through simulation and also real life experiments that game participants find it extremely difficult to perform well in this game. Their decisions commonly result in large divergences which are far from optimal behaviour. These result in large oscillations, deterministic chaos and other forms of complex behaviour (Mosekilde and Larsen, 1988).

In order to identify the most optimal strategies for all game participants, numerous optimisation techniques have been successfully applied in the BDG, such as Genetic Algorithms (GAs) and Particle Swarm Optimisation (PSO) (Strozzi et al., 2007; de Souza et al., 2000). Most existing research has focused on one common assumption that all the game participants have one common goal which is to minimise the cost of the whole supply chain. However, in the real world, it is intuitive to argue that in the BDG that any two participants are only ever concerned about their own performance in the game. A retailer is never concerned about the performance of the factory, and vice-versa. Thereby, due to the conflicting preferences and incomplete information involved in these interactions the reality is that participants determine their actions based on their own circumstances. This reflects their inherent bounded rationality (Simon, 1997) which is a fundamental

factor in real world supply networks. In this paper, we aim to address this issue and provide a series of simulations investigating this property.

In this paper, we must consider game participants as agents which have individual strategies. In order to examine the effects of bounded rationality on the evolution of strategies in the BDG, we require a evolutionary framework which offers a means of evolving these agent strategies independently in each sector of the BDG. Therefore, we propose a coevolutionary architecture to address this issue. Coevolution is a process of mutual adaptation that occurs amongst a set of individuals that interact strategically in some domain. Each agent's strategy is evaluated based on its fitness, and subsequently this is used to determine the evolution of strategies over successive generations. This adopts the same approach as a genetic algorithm, however, we consider each sector of the BDG to be a distinct agent population upon which a genetic algorithm is applied. Crossover and mutation only occur within a given sector and not between sectors of the game. This is therefore considered the coevolution of agent strategies. Strategies across all sectors of the game evolve in parallel, and influence each other only through their game interactions.

In this paper we are particularly interested in the effects of greater agent autonomy of game participants. Therefore, we seek to explore the effects of allowing these individuals determine their strategies based on their own individual choices and preferences. We propose allowing these individuals evaluate their strategies based on individual fitness functions whereby their sole consideration is the minimisation of their costs in the game. This allows us to explore the effects of individuals evolving their strategies in a more bounded context. Therefore, their decisions are determined through their limited view of the world around them and their goal to maximise their own benefit, thereby reflecting a greater degree of bounded rationality.

In contrast to evaluating fitness individually, we will also show a series of simulations which evaluate each participants fitness on the overall performance of the supply chain and not their own costs. Thus, reflecting a more traditional approach commonly used when attempting to optimise the BDG. Thereby showing two sets of experiments which reflect the evolution of agent strategies in the BDG. Two evolved BDG strategy sets will be discussed:

1. Strategies evolved when individual fitnesses are evaluated based on their individual costs. This represents an approach that each agent chooses to minimise their own cost regardless of the whole supply chain cost in the BDG. Thereby this reflects agents who are individually rationality.
2. Strategies evolved when individual fitnesses are evaluated based on collective supply chain costs. In the BDG, this represents an approach where all agents are rewarded for making choices which maybe individually irrational for the sake of the entire supply chain. This correlates to a collectively rational approach and thus reflects agents which behave with group or collective rationality.

Throughout the simulations presented in this paper we will compare the effects of using these two approaches in a coevolutionary framework. Agent strategies will evolve using genetic algorithms independently in each sector of the game. Thereby this reflects that each sector of the game has certain unique properties related to its position relative to the other sectors. The agent populations in each sector will co-evolve independently over many generations.

This differs significantly to existing approaches to BDG optimisation which focus on optimising strategies based on a global fitness function in order to minimise the costs of the entire supply chain. This paper will address a number of important research questions:

1. What are the effects on the BDG when agents use individual rationality or group rationality?
2. How do the most fit strategies for each sector from the BDG evolve over time?
3. In light of our results what are the broader implications for BDG optimisation and supply chains?

These research questions will be referred to regularly throughout this paper and answered directly in the Conclusions section. The following sections of this paper are structured as follows. In Section 2, we will discuss background research and in Section 3 we will outline our simulator design. In section 4 we will outline our experimental setup. Section 5 will provide a detailed examination of our experimental results. In Section 6 we will outline our conclusions, while finally in Section 7 we will briefly summarise the contributions of this paper and outline some future work.

2 Background Research

2.1 Beer Distribution Game

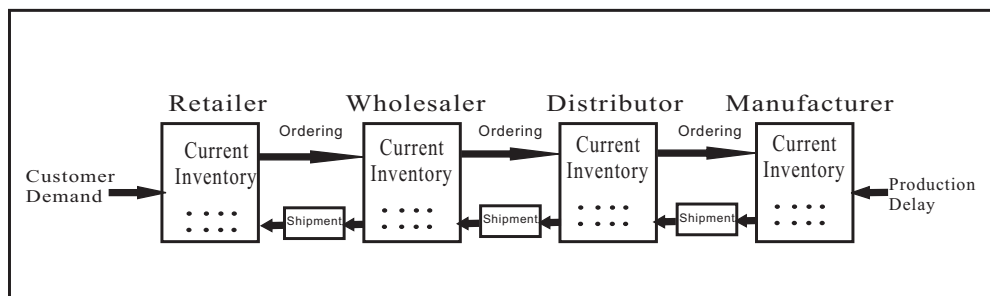


Figure 1: The structure of the Beer Distribution Game

The BDG is a classic supply chain game which has been widely used in the domain of system dynamics and supply chain management (Sterman, 1987; Huang et al., 2003; von Lanzener et al., 2001; de Souza et al., 2000). This game offers a simplified implementation of common real world production and distribution systems. As shown in Fig. 1, this system consists of four participants: Retailer, Distributor, Wholesaler and Manufacturer (R, W, D, M). Each participant has control and responsibility for its own inventory.

2.2 The Beer Distribution Game Equations

In this section, we briefly outline the decision rules for stock management in each sector of the BDG. More detailed information can be found in Sterman (1987).

The decision rule for each sector utilizes information locally available to each sector. The information includes the current demand from the direct downstream sector, the current inventory level and supply line level. Each sector does not know the other sectors' inventory level or supply line information. This incomplete information is what makes the BDG notoriously difficult. In a given week t , each sector applies this rule in order to satisfy its expected demand Le_t , reduce the discrepancy between the desired and actual stock AS_t , and maintain an adequate supply line of unfilled orders ASL_t . First, orders O_t must be nonnegative:

$$O_t = \max(0, IO_t^*) \quad (1)$$

The indicated order IO_t^* is represented as follows:

$$IO_t^* = Le_t + AS_t + ASL_t \quad (2)$$

The expected demand Le_t is weighted between the actual demand L_{t-1} and the expected demand at week $t - 1$:

$$Le_t = \theta * L_{t-1} + (1 - \theta) * Le_{t-1}, \quad (3)$$

In Equation (3), the parameter θ is the weighting towards considering the actual demand in the last week.

The adjustment of the stock level AS_t has a linear relationship to the discrepancy between the desired stock S^* and the actual stock S_t at week t . This is formulated as follow:

$$AS_t = \alpha(S^* - S_t) \quad (4)$$

The stock adjustment parameter α is the fraction of the discrepancy ordered each period. This is usually represented in the range $(0 \leq \alpha \leq 1)$.

The adjustment for the supply line is formulated analogously as

$$ASL_t = \alpha_{SL}(SL^* - SL_t) \quad (5)$$

where SL^* is the desired supply line and SL_t is the actual supply line, and α_{SL} is the fractional adjustment rate for the supply line.

Defining $\beta = \alpha_{SL}/\alpha$ and $S' = S^* + \beta SL^*$ and then we can get:

$$IO_t^* = Le_t + \alpha(S' - S_t - \beta SL_t) \quad (6)$$

where β is the fraction of the supply line taken into account by the sectors. This parameter is usually represented in the range $(0 \leq \beta \leq 1)$. If $\beta = 1$, the participant factors in all orders in the supply line or conversely, if $\beta = 0$, the participant factors in no orders in the supply line.

The combination of the adjustment parameters (α, β) corresponds to a set of strategies for the game participants in their inventory management. In this paper, we will allow each participant to

freely choose their strategies and evolve with each other within our coevolutionary model.

2.3 The Objective Function

The objective of each sector is to minimise cumulative costs over N weeks by keeping inventories as low as possible while avoiding out-of-inventory conditions which cause backlogs. The BDG commonly uses the following costs to penalise inventory holding and backlogs. The cost of inventory holding is \$0.5 for each case of beer per week and the cost of backlogs is \$2.0 for each case of beer per week. It is intuitive for a player to order more beer when inventory falls below a desired level. Similarly a player is likely to order less beer when stocks begin to accumulate.

Table 1 shows the objective functions for individuals in each sector of the game. These are the retailer C_R , wholesaler C_W , Distributor C_D , and manufacturer C_M . Finally we show the collective costs C_T of all sectors over N weeks. In this table, the inventory level and backlogs at each end of week are presented as INV and BL , respectively.

Table 1: The Objective Function

	Objective Function
Retailer	$C_R = \sum_{i=1}^N (0.5 * INV_R^i + 2.0 * BL_R^i)$
Wholesaler	$C_W = \sum_{i=1}^N (0.5 * INV_W^i + 2.0 * BL_W^i)$
Distributor	$C_D = \sum_{i=1}^N (0.5 * INV_D^i + 2.0 * BL_D^i)$
Manufacturer	$C_M = \sum_{i=1}^N (0.5 * INV_M^i + 2.0 * BL_M^i)$
Whole Supply Chain	$C_T = C_R + C_W + C_D + C_M$

2.4 Coevolutionary Approaches

In this section, we introduce a number of alternative evolutionary approaches and coevolutionary approaches. Evolutionary approaches (EAs) are generic population-based metaheuristic optimization algorithms such as Genetic Algorithms (GAs). EAs are inspired by biological evolution: reproduction, recombination, mutation, and selection. In general, each individual in a population represents a solution to a optimization problem in some domain. The individual's fitness symbolizes the quality of the solution and subsequently plays a vital role in its evolution. EAs share a similar framework which involve firstly randomly generating a population and subsequently evaluating each individual. Then the selection, recombination, mutation operators will be used. Over repeated iterations of the specified process, a population will evolve with respect to the fitness landscape involved. Alleles associated with the most fit individuals are propagated throughout the population while less fit alleles are less likely to propagate. EAs have been successfully applied to optimisation problems such as the traveling salesman problem, scheduling, adaptive control, supply chain management, etc (Goldberg and Lingle, 1985; Michalewicz, 1992; Sourirajan et al., 2009).

Coevolutionary approaches are also a subdomain of EAs. The main difference is that coevolutionary approaches usually involve many populations. Each population contains part of a solution

compared with the standard EAs. Coevolutionary approaches also use the recombination, mutation, selection operators. Coevolutionary approaches have been used extensively in a number of difficult problem domains such as function optimisation (Hillis, 1990; Potter and Jong, 1994; Potter and De Jong, 2000; Yang et al., 2008) and inventory control optimisation (Eriksson and Olsson, 1997).

3 Simulator Design

In this paper, we use a co-evolutionary approach to investigate individual rationality and group rationality in the BDG. Each agent must make decisions according to their specified goal. In this section, we will first outline our coevolutionary framework, and then our two alternative agent fitness evaluation criteria.

3.1 Coevolutionary Framework

In this section, we will present our coevolutionary algorithm design, and then present our simulation design as a whole.

3.1.1 Coevolutionary Algorithm Design

In these simulations we use a coevolutionary algorithm to allow each agent population to evolve independently in each sector of the BDG. The Figure 2, shows a pseudo code representation of our coevolutionary approach. As shown, we first randomly generate M populations where each population has N individuals. This results in a uniform distribution of agent strategies in each agent population. Individual agents from each sector of the game then participate in instances of the BDG. This provides a means of determining their relative performance in the game and allows us to allocate them fitness values. Subsequently, the selection, crossover and mutation operators will be applied in each population independently. This process will continue repeatedly until it reaches the max allowed generations.

```

Initialise  $M$  populations and each population has  $N$  individuals;
for ( $i=0; i<MAX\_ALLOWED\_GENERATIONS; i++$ ) {
  for(each population) {
    for(each individual ( $a$ )) {
      Calculate fitness =  $f(a)$ ;
    }
    Selection;
    Crossover;
    Mutation;
  }
}

```

Figure 2: Our Coevolutionary Algorithm




Figure 3: The Simulation Model

3.1.2 The Simulation Model

The Figure 3 shows the entire simulation model. In this model, there are four agent populations: retailer P_R , wholesaler P_W , distributor P_D and manufacturer P_M .

Through coevolution each population evolves independently from the other agent populations. Figure 3 indicates the particular phase of the simulation model when we are evaluating the agents from the retailer agent population P_R . Each of the other agent populations P_W, P_D, P_M provides agent representatives to participate in the process. As a result, each agent from the retailer will play an instance of the BDG with all the representative agents from the other sectors. This means that each agent from the Retailer Population (R) will participate in N^{M-1} interactions with the other representative agents. The representative agents will fulfill only the role defined by their associated sector. For example, an agent from the distributor agent population will always play the role of a distributor in all games it participates in. This is similar to a round robin type implementation whereby each of the relevant agents are paired in order to participate in a instance of the BDG.

Furthermore, in order to evaluate an agent's fitness, we will average their total cost over all these interactions (N^{M-1}) and then assign this value as their individual fitness. We believe this is the fairest approach in order to compare the performance of individuals in each agent population. However, our approach is very computationally expensive ¹.

After evaluating all the agents from the retailer agent population, we will apply this process to the other agent populations in a sequential pattern. The sequential pattern is: retailer, wholesaler, distributor and finally manufacturer. This is similar to other coevolutionary approaches such as Potter and Jong (1994). However, this process could also be evaluated purely in parallel. This reflects the process of evaluating each agent's fitness in the specific sector population concerned. Once all agent's in the sector population have an associated fitness we then apply selection, crossover,

¹Each generation involves N^M evaluations or the BDG simulations.

mutation operators to that sectors agent population as shown in Figure 2.

3.2 Individual Versus Group Rationality

In the BDG, it is desirable for all four sectors to fully cooperate with each other in order to reduce the whole supply chain cost. However, in reality the behaviour of human participants in the BDG has shown that this is difficult to achieve and rarely occurs. There are at least two reasons. Firstly, each sector does not have complete information to manage its inventory. For example, all the sectors do not know the actual customer demand except for the retailer. Secondly, real world individuals are mostly self-interested, and are primarily concerned with maximising their own performance. This reflects that there are conflicting preferences between the sectors in the BDG. Thus, as in the real world, peoples behaviors are effected by bounded rationality as they have a limited view of the world around them. Therefore, they simply make decisions in order to maximise their own benefit at any moment in time.

As we have mentioned previously, this paper aims to investigate the effects of bounded rationality in the BDG. We use two fitness criteria as a means of modeling this feature. Two alternative fitness criteria are specified as follows:

1. **Individual Fitness Function:** The agents from each population P_R , P_W , P_D , and P_M use C_R , C_W , C_D and C_M respectively as shown in Table 1. These objective functions reflect that each agent seeks to minimise its individual cost regardless of the whole supply chain cost. This reflects how these agents are rewarded in our evolutionary algorithm based on their individual performance, thereby ignoring the performance of other individuals in the supply chain. This has the effect of rewarding individually rational behavior even if this behavior may not be collectively rational for the entire supply chain. The agents adjust their strategies according to their own self-interest and promote their own fitness criteria over each generation in our coevolutionary model. In the following, we will refer to this approach as *IFF*.
2. **Global Fitness Function:** All the agents from P_R , P_W , P_D , and P_M use the whole supply chain cost as their objective function C_T in Table 1. This indicates a situation whereby all agents are rewarded through our evolutionary algorithm based on the performance of the entire supply chain. Thereby ignoring whether their actions maybe individually irrational. They are rewarded if they help maximise the efficiency of the overall supply chain, even if this is at their own individual expense. Thereby, compared with *IFF*, this represents a reduced degree of bounded rationality with respect to the BDG. In the following discussion, we will refer to this approach as *GFF*.

These two fitness evaluation approaches provide us with a basis to compare the differences and similarities between the effects of the bounded rationality in the BDG when individuals behave individual rationality *IFF*, or alternatively using their collective or group rationality *GFF*.

4 Experimental Setup

In this section, we will outline the parameters involved in our simulation model and the BDG. In our coevolutionary model we must simulate the BDG a number of times. The following parameters

are used throughout our simulations: The total simulation length is 50 weeks. The customer demand is initially four cases per week and increases to eight cases per week in week 5 and remains at that level thereafter. The α and β values are all in the range of $[0, 1]$. The θ value used is 0.25 (see Equation (3)). Delays equal 4 weeks and the parameter S' in Equation 6 is set to 17.

As shown in Figure 3, there are four sectors representing four agent populations in the simulation model. Each population size is set to $N = 10$ and the total number of iterations (generations) is $= 150$. The population size used is quite small due to the computationally intensive nature of this approach. A selection operator which selects the best individual (as determined by their fitness) is used and a selection rate of 0.9 is applied. The crossover rate is 0.85 and the mutation rate is 0.2.

5 Experimental Results

In this section, we will outline a series of experimental results from our simulations. Firstly, we will examine the evolution of the most fit strategies and their effects on agent fitnesses in each sector of the BDG. Subsequently, we will show the overall effects on the supply chain when *IFF* and *GFF* are used. We will compare both approaches and outline their differences and similarities.

Please note that each agent records a number of important metrics in our simulation. Firstly, each agent records the average cost incurred by the whole supply chain (C_T) over a number of interactions (N^{M-1}) in a generation. Secondly, each agent records its own average cost incurred over N^{M-1} interactions. Thirdly, each agent has its own individual strategies (α and β). Finally, all agents record their own fitness in each generation. As mentioned above, the fitness value equals to the average of its individual cost or the average of the whole supply chain cost depending on using *IFF* or *GFF*. Our following experimental results are from the agents that have the best fitness among each sector agent population.

5.1 Strategy Evolution

In this section, we will show a series results showing the evolution of the most fit strategies and the best fitnesses for each sector from each respective agent population. We will show the differences and similarities between the results from the agents use *IFF* and *GFF*.

5.1.1 Agent Strategies

Each agent has two parameters α and β which represent its strategies. A high α represents an agent's high attention to the inventory, a high β represents an agents high attention to the supply line comparison with the inventory. Existing research has shown that the optimal strategies for each sector is that α and β are close to 1 (Sterman, 1987).

Figure 7 shows best strategies α value for each sector changes over generation when each agent uses *IFF*. Figure 8 represents the best strategies when all agents use *GFF*. Figure 9 shows the β value for the best strategies in each sector over successive generations when using *IFF*. Figure 10 shows the same data when all agents use *GFF*. Each set of results shows the strategy changes over one typical run while a corresponding figure shows the average changes over 50 runs.

There are several features from these experiments. Firstly, we observe that there are similar trends between the data in each single run and the corresponding averaged data in Figures 7,8,9,10.

Secondly, the (α and β) values of the best strategies in each sector increase at quite rapidly initially and then remain quite steady over successive generations. This reflects that individuals pay more attention to their inventory and supply line and obtain better performances over repeated generations. This performance is shown in Figure 4, 5. Finally, our results show the difficulty for individuals using both *IFF* and *GFF* to evolve high α and β values. This is an indication of the complexity of the fitness landscape inherent in this game environment.

5.1.2 BDG Sector Costs

Figure 4 shows each associated sector cost from the best agents over successive generations agent's use *IFF*. Figure 5 shows the same data when each agent uses *GFF*. We outline the changes in one typical run while also showing data showing average changes over 50 runs.


Figures 4 and 5 show no major differences between our typical run and the averaged data. This indicates the relative stability of the overall evolved behaviors. We notice that cost falls over time as our populations converge. This reflects the ability of our coevolutionary framework to promote the best performing strategies in each of our agent populations. For the BDG, this process reflects how agent strategies evolve over successive generations to converge on the most fit solutions that can be identified in that particular game environment. This process is guided by each agent's fitness criteria *IFF* and *GFF*. Finally, we observe significant differences in how agents perform in Figure 4. However, no significant differences are apparent in Figure 5. More specifically, in Figure 4 the retailer always performs best, with the wholesaler, distributor, and manufacturer following in that order of performance. This indicates the direct implications of allowing each of these sectors evolve independently through *IFF*. As a result these agent populations are continuously tuning their strategies to perform as well as possible without regard for the entire supply chain. The opposite is identifiable in the results outlined for *GFF* as shown in Figure 5.

5.1.3 BDG Supply Chain Costs


In this section, we will present how the whole supply chain cost evolves over successive generations. Figure 6 shows the average whole supply chain cost as our populations evolve over a number of generations. Figure 6(a) shows the changes from one typical run while Figure 6(b) shows these same changes averaged across 50 runs. We observe from Figure 6 that overall costs are higher for the entire supply chain when *IFF* is used. These results are consistent throughout numerous experimental runs as we can see from the data. Despite both approaches reducing overall supply chain costs, it is clear that *GFF* is far more successful at this task. This data confirms our belief that when agents act solely to minimise their own costs they will in effect end up undermining their own performance and that of the entire supply chain. Instead they must act counter intuitively and make initial sacrifices for the sake of the entire supply chain. This displays the clear impact and significance of individual and group rationality in these environments.

5.2 Supply Chain Analysis

In this section we will present a series of results showing the performance of the agent populations from each sector while participating in the BDG. We will show the differences and similarities between the results from our simulations involving *IFF* and *GFF*.




(a) Average Sector Costs (1 Run)




(b) Average Sector Costs (50 Runs)

Figure 4: Average Sector Costs (1 Run)Vs(50 Runs) Using *IFF*

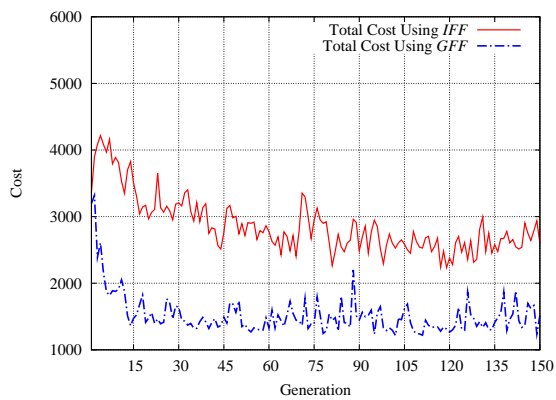


(a) Average Sector Costs (1 Run)

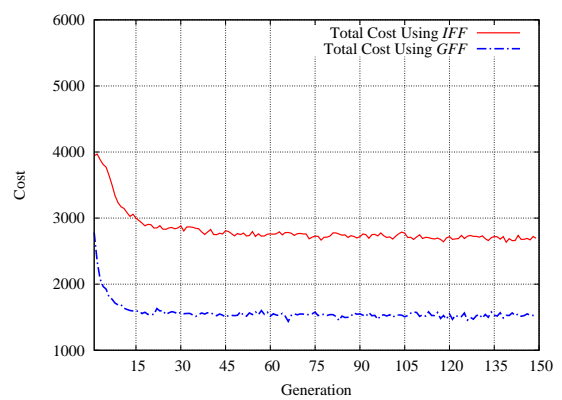


(b) Average Sector Costs (50 Runs)

Figure 5: Average Sector Costs (1 Run)Vs(50 Runs) Using *GFF*




(a) Average Total Cost (1 Run)




(b) Average Total Cost (50 Runs)

Figure 6: Average Total Cost (1 Run)Vs(50 Runs) using *IFF* and *GFF*




(a) α (1 Run)




(b) Average α (50 Runs)

Figure 7: α (1 Run)Vs(50 Runs) using *IFF*




(a) α (1 Run)




(b) Average α (50 Runs)

Figure 8: α (1 Run)Vs(50 Runs) using *GFF*



(a) β (1 Run)



(b) Average β (50 Runs)

Figure 9: β (1 Run)Vs(50 Runs) using *IFF*




Figure 10: β (1 Run)Vs(50 Runs) using *GFF*

5.2.1 Simulation Results using Individual Rationality (*IFF*)

In this section we will present our experimental results from our evolved agent populations when using *IFF* as the fitness evaluation function.

Table 2: Results from Individual Rationality (*IFF*)

	Worst	Best	μ	σ
C_R	166.1	116.7	136.6	13.4
C_W	387.0	253.9	303.0	27.9
C_D	1009.7	679.1	847.2	82.1
C_M	1538.9	998.1	1262.7	100.6
C_T	3177.1	2123.3	2618.2	215.0

The results presented in Table 2 are from 50 individual runs of our coevolutionary algorithms when all agents use *IFF*. Thereby, this reflects that each agent strategy evolves based on their own costs incurred. Therefore we show the most relevant data in bold italic font. The data shows the cost associated with each potential role in the BDG. We also show the average total supply chain cost T_C incurred by the four sector populations in the BDG. The data in Table 2 shows the maximum cost (Worst), the minimum cost (Best), the average cost (μ) and standard deviation (σ). From this table, we observe that the more upstream individuals in the BDG incur higher costs and σ . The underlying reason for this is that the knowledge of the downstream orders tends to be distorted and can subsequently misguide the upstream individuals in their inventory and production decisions. This distortion tends to increase and thus causes a larger inventory fluctuation as one moves upstream. This is a phenomenon commonly referred to as the “bullwhip effect” (Lee et al., 1997).

5.2.2 Simulation Results using Collective Rationality (*GFF*)

In this section we will present our experimental results from our evolved agent populations when using *GFF* as the fitness evaluation function. As in the previous experiment described in section 5.2.1, the data presented here in Table 3 is from 50 experimental runs. We will discuss these results in detail when comparing them to those evolved using *IFF* in the following section.

Table 3: Results from Collective Rationality (*GFF*)

	Worst	Best	μ	σ
C_R	636.5	216.1	429.0	107.0
C_W	564.8	157.2	326.4	82.7
C_D	582.8	184.2	297.5	93.6
C_M	706.8	237.5	387.7	108.2
C_T	1843.5	1103.7	1406.9	206.8

5.2.3 Comparison of Individual Rationality (*IFF*) and Collective Rationality (*GFF*)

In this section we will identify the differences and similarities between our two evolutionary models *IFF* and *GFF*.

Table 4: The Percentage Cost Reduced When The Agents Use Individual Rationality (*IFF*) Compared with That The Agents Use Collective Rationality (*GFF*)

	Worst	Best	μ
C_R	73.9%	46.0%	68.2%
C_W	31.5%	-61.5%	7.2%
C_D	-73.2%	-268.7%	-184.8%
C_M	-117.7%	-320.3%	-225.7%
C_T	-72.3%	-92.4%	-86.1%

We also examined our evolved agent strategies when *GFF* is used to evaluate their fitnesses. From the data shown in Table 3 we can identify clear differences with the results shown in Table 2. We observe that retailer is the worst performer, as identified by μ in this table. This primarily due to individuals throughout the supply chain adjusting their strategies based on the common goal of minimising the whole supply chain cost. Individuals will sacrifice their individual profit in order to benefit the performance of the overall supply chain.

The results presented in Table 4 show the cost reduction when the agents use *GFF* compared with that the agents use *IFF*. The data in this table shows the performance differences and the effects on the whole supply chain between individual rationality and collective rationality. The agents using *IFF* had attempted to reduce their own cost, however the benefits of this falls significantly as one moves upstream. From this table, we can see only the retailer agents gains from this

approach while the distributor and manufacturer suffer significant losses. The wholesaler performs quite similarly in both *IFF* and *GFF*. These observations reflect many of the characteristics we would expect given the existing research involving supply chains. If downstream sectors manage their inventory selfishly, this will impact significantly on any upstream sectors. These simulation results provide further evidence of the importance of cooperation and information between the sectors of the supply chain.

Table 5: The statistical significance using *IFF* vs *GFF*

<i>T</i> test	C_R	C_W	C_D	C_M	C_T
<i>P</i>	0.00%	6.28%	0.00%	0.00%	0.00%
<i>T</i> Value	-18.698	-1.896	31.219	41.878	28.712
Significance	TRUE	FALSE	TRUE	TRUE	TRUE

To further reinforce our observations regarding the differences between individuals evolved using individual rationality and collective rationality, we conducted a statistical *T* test to demonstrate the significance of the differences between our two sets of results (*IFF* and *GFF*). Furthermore, we use the conventional criteria to determine whether differences are significant (S). That is if the two tailed p value is less than 5%, the difference is statistically significant (TRUE), or else, it is not (FALSE). The data in Table 5 verifies our earlier analysis, and that the behaviours evolved in the two models are significantly different. The only sector not showing a significant statistical difference was the Wholesaler and we believe this is purely down to his location in the supply chain relative to its peers. This results is expected and confirms that the results of our two models are significantly different.

6 Conclusions

This paper presents a new and alternative approach to evolving strategies for the BDG. This includes a co-evolutionary architecture which provides a means of evolving strategies across the various sectors of the BDG. We propose two alternative sector evolution preference, one reflects individual rationality and the other reflects group or collective rationality. A comprehensive definition and comparison between these preferences has been presented.

Earlier in this paper, we posed a number of important research questions. In response to our first research questions, it is clear that in most cases individual and collective performance improves significantly when individuals avoid acting purely through their individual rationality but instead determine their behavior based on group rationality. The effects on upstream individuals are more significant than those downstream. Only the downstream retailer gains when individuals use their individual rationality. The upstream sectors suffer heavily in this case, and this thereby reflects aspects of the “bullwhip effect”.

Our second research question refers to the changes of the most fit strategies for each sector over successive generations in our simulation model. We observe that the most fit strategies for each sector tend to increase over time. However, these still fail to reach optimal strategies in our simulation model. This stems from each sectors in the BDG suffering through their lack of information.

In the BDG, each individual makes its decisions based on local information such as the demand from their direct downstream sector, the current inventory level, and supply line level. They are never aware of the actual custom demand (except for the retailer) or their neighbor's inventories levels. Furthermore, we use a small number of agents in each population in our coevolutionary model. This would undermine our ability to evolve certain strategies. In future, we would like to investigate this issue further.

Our results have shown that bounded rationality is a significant factor when examining the BDG. The conflict between individually rational choices and collectively rational choices are very difficult for BDG participants to clearly identify. Our results also have demonstrated the difficulty of trying to make individually rational choices in the BDG, as this can directly contribute to poorer performances for the individual concerned and the entire supply chain. This reflects many of the properties of a typical social dilemma. In summary, our results underscore the importance of cooperation in supply chain management.

7 Summary and Future Research

This paper has presented a novel extension to existing research involving the BDG. This includes simulating BDG participants which have greater autonomy to determine their strategies based on their individual preferences. Furthermore, we introduce for the first time a coevolutionary model to examine the effects of bounded rationality on the sectors in the BDG. However, a number of fundamental factors influence this study. Firstly, the degree of bounded rationality reflected by the game participants. As we have shown this has significant implications for performance of any given supply chain. Even in the limited case of this paper, we have seen the effects of individual rationality and group rationality on the supply chain. A number of limitations are inherent in our simulations. Firstly, our study only involves one very simple customer demand pattern and also includes very small agent populations. We hope to address these limitations in future research.

8 Acknowledgment

The authors would like to gratefully acknowledge the continued support of Science Foundation Ireland.

References

- de Souza, R. et al. (2000). Supply chain dynamics and optimization. *Integrated Manufacturing Systems*, 348–364.
- Eriksson, R. and B. Olsson (1997). Cooperative coevolution in inventory control optimisation. In *Proceedings of the Third International Conference on Artificial Neural Networks and Genetic Algorithms*, University of East Anglia, pp. 583–587. Springer.
- Forrester, J. W. (1961). *Industrial Dynamics*. Cambridge, MA: MIT Press.

- Goldberg, D. E. and R. Lingle, Jr. (1985). Alleles and the traveling salesman problem. In *Proceedings of the 1st International Conference on Genetic Algorithms*, Hillsdale, NJ, USA, pp. 154–159. L. Erlbaum Associates Inc.
- Hillis, W. D. (1990). Co-evolving parasites improve simulated evolution as an optimization procedure. *Phys. D* 42(1-3), 228–234.
- Huang, G. Q. et al. (2003). The impacts of sharing production information on supply chain dynamics: a review of the literature. *International Journal of Production Research* 41(7), 1483–1517.
- Knolmayer, G., R. Schmidt, and S. Rihs (2007). Teaching supply chain dynamics beyond the beer game. In *Institute of Information Systems University of Bern*.
- Lee, H. L., V. Padmanabhan, and S. Whang (1997). Information distortion in a supply chain: the bullwhip effect. *Manage. Sci.* 43(4), 546–558.
- Michalewicz, Z. (1992). *Genetic algorithms + data structures = evolution programs*. Artificial Intelligence. Berlin etc.: Springer-Verlag. XIV, 250 p. .
- Mosekilde, E. and E. R. Larsen (1988). Deterministic chaos in the beer production-distribution model. In *System Dynamics Review*, Volume 4.
- Potter, M. A. and K. A. De Jong (2000). Cooperative coevolution: An architecture for evolving coadapted subcomponents. *Evol. Comput.* 8(1), 1–29.
- Potter, M. A. and K. A. D. Jong (1994). A cooperative coevolutionary approach to function optimization. In *PPSN III: Proceedings of the International Conference on Evolutionary Computation. The Third Conference on Parallel Problem Solving from Nature*, London, UK, pp. 249–257. Springer-Verlag.
- Simon, H. A. (1997). *Models of Bounded Rationality - Vol. 3 Empirically Grounded Economic Reason*. MIT Press.
- Sourirajan, K., L. Ozsen, and R. Uzsoy (2009). A genetic algorithm for a single product network design model with lead time and safety stock considerations. *European Journal of Operational Research* 197(2), 599 – 608.
- Sterman, J. (1984). Instructions for running the beer distribution game. *System Dynamics Group Working Paper D-3679, MIT, Sloan School of Management, Cambridge, MA*.
- Sterman, J. (1987). Modeling managerial behavior–misperceptions of feedback in a dynamic decisionmaking experiment. Technical report.
- Strozzi, F. et al. (2007). Beer game order policy optimization under changing customer demand. *Decis. Support Syst.* 42(4), 2153–2163.
- Thomsen, J. S., E. Mosekilde, and J. D. Sterman (1992). Hyperchaotic phenomena in dynamic decision making. *Systems Analysis and Modelling Simulation* 9, 137–156.

von Lanzener, C. H. et al. (2001). Coordinating supply chain decisions: an optimization model. *OR Spectrum* 24(1), 59–78.

Yang, Z., K. Tang, and X. Yao (2008, August). Large scale evolutionary optimization using cooperative coevolution. *Information Sciences* 178(15), 2985–2999.