Multi-Agent Simulation of Collaborative Strategies in a Supply Chain

Thierry Moyaux, Brahim Chaib-draa
FOR@C, CENTOR & DAMAS
Université Laval - Pavillon Pouliot
Département d'Informatique et de Génie Logiciel
Sainte-Foy G1K 7P4 (Québec, Canada)
{moyaux, chaib}@iad.ift.ulaval.ca

Sophie D'Amours
FOR@C & CENTOR
Université Laval - Pavillon Pouliot
Département de Génie Mécanique
Sainte-Foy G1K 7P4 (Québec, Canada)
sophie.damours@gmc.ulaval.ca

Abstract

The bullwhip effect is the amplification of the order variability in a supply chain. This phenomenon causes important financial cost due to higher inventory levels and agility reduction. In this paper, we study, for each company in a supply chain, the individual incentive to collaborate to reduce this problem. To achieve this, we simulate a supply chain inspired by the Québec forest industry, in which each company is an agent that uses one of three ordering schemes. Each ordering scheme represents a level of collaboration. One run of the simulation is done with fifty (50) weeks for each of the $3^6 = 729$ *combinations of these 3 ordering schemes among the 6 companies of the simulation. In each run, we evaluate each company's inventory holding and backorder costs. These outcomes are used to build a game in the normal form, which is next analyzed using Game Theory. In particular, we find two Nash equilibria incurring the minimum cost of the supply chain. We also note that there are no Nash equilibria in which some companies do not collaborate: collaborating companies have no incentive to stop collaboration.*

1. Introduction

A supply chain is a set of autonomous business units producing and distributing products. The bullwhip effect is the amplification of the order variability in such a supply chain. In other words, the bullwhip effect is a distortion of demand information when this information is transmitted as orders along the supply chain down to the *n* tier suppliers. Therefore, this deformation of information does not only interest Supply Chain Management, but also Computer Science which is the study of information processing.

The bullwhip effect costs money due to higher inventory levels and supply chain agility reduction. It is a problem of coordination and collaboration between selfish agents (i.e., autonomous business units of the supply chain). In fact, downstream companies (e.g. retailers) do not suffer directly from it while they are in position to reduce it, as the bullwhip effect amplifies along the supply chain, whereas upstream companies (i.e. raw material suppliers, such as a forest) suffer from this phenomenon while they cannot counter it.

Solving this problem requires collaboration in the supply chain, but the problem lies in the fact that companies have conflicting objectives, and thus, may not have an incentive to collaborate. Information sharing, here demand information transmission, is one method of collaboration which is often said to be the solution to the bullwhip effect [4, 6, 14]. We have previously designed a decentralized decision process (i.e., an ordering scheme) based on information sharing to minimize the bullwhip effect without neglecting the importance of inventory management and operational constraints [10, 11]. Up to now, we have only examined the efficiency of these schemes in multi-agent simulations of a homogeneous supply chain, i.e., a supply chain where all company-agents use the same ordering scheme. In this paper, we consider a heterogeneous supply chain by letting company-agents use different ordering schemes. This allows us to study companies' incentives for using our decision process and therefore for collaborating. We describe this problem of incentive for collaborating in Section 2.

We use concepts from Game Theory to analyze the simulation outcomes. Precisely, outcomes are used to build a game in the normal form, and we look in this game for strictly dominated strategies, Nash equilibria and minimum supply chain costs. A strictly dominated strategy is a company's scheme that is never used because it always incurs the highest cost for the company no matter what is chosen by the rest of the supply chain. A Nash equilibrium is a situation where no company has an incentive to change its ordering scheme. The minimum overall cost is the best situation for the whole supply chain, even if some companies are not gaining as much as they could. These Game Theory concepts are presented in Section 3.

Multi-agent simulations are carried out to build the studied games. That is, each company is seen as a reactive agent that applies a given ordering scheme, and we simulate the corresponding supply chain in order to evaluate inventory and backorder costs for each company. In each run of the simulation, each company-agent applies one of three ordering schemes. We restrict the number of possible schemes to 3, because the number of combinations of schemes among the 6 companies in the simulation is $3^6 = 729$. For each combination, a run of the simulation is carried out and the cost incurred by each company is written in a $3 \times 3 \times 3 \times 3 \times 3 \times 3$ matrix. This matrix is a game in the normal form that we analyze with the above mentioned game-theoretic concepts. The simulation model, the three ordering schemes and the analysis of the simulation outcomes are given in Section 4

This analysis shows that two Nash equilibria incur the minimum supply chain cost. Intuitively, this minimum should be incurred when every company fully collaborates, but surprisingly, there is one company that only collaborates partially in these two equilibria, while all other companies fully collaborate. The supply chain should therefore try to reach one of these two equilibria, because they reduce the overall supply chain cost while no company would be better off leaving them. This interpretation of the simulation outcomes is developped in Section 5.

2. Problem Statement and Motivation

We now present the bullwhip effect and then, a solution that we have previously proposed to reduce it. Next, we outline the motivation for the approach presented in this paper: we would like to know if selfish companies have individual incentive to use such a solution for reducing the bullwhip effect.

2.1. The Bullwhip Effect

The bullwhip effect is a deformation of demand information when this information is transmitted as orders move down the supply chain. Figure 1 shows how the bullwhip effect propagates in a simple supply chain with only three companies in the specific case of the forest supply chain: a retailer, a wholesaler and a pulp mill. The retailer sells to the customer and buys from the wholesaler, the wholesaler sells to the retailer and buys from the pulp mill and the pulp mill sells to the wholesaler and buys from an unknown supplier. The ordering patterns of the three companies "share a common, recurring theme: the variabilities of an upstream site are always greater than those of the downstream site" [6]. As a variability, the bullwhip effect is measured by the standard deviation σ of orders (note that means μ of orders are all equal in our example).

There are several consequences of the bullwhip effect: this effect incurs costs due to higher inventory levels and supply chain agility reduction, decrease of customer service levels, ineffective transportation, missed production schedules... In fact, such fluctuations of the demand lead every participant in the supply chain to "stockpile because of a high degree of demand uncertainties and variabilities" [7]. An insight of the importance of this problem is given by Carlsson and Fullér [3] who estimate that the costs incurred by the bullwhip effect are 200-300 MFIM (40-60 millions USD) annually for a 300 kton North-European paper mill.

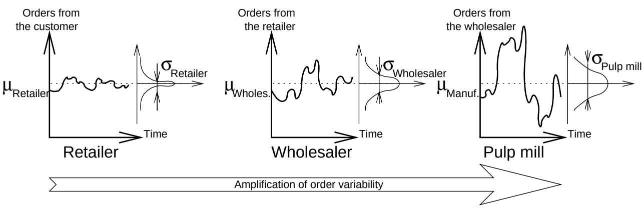

Figure 1. The bullwhip effect [6, 7].

2.2. Information Sharing as a Solution to the Bullwhip Effect

We have previously proposed a decentralized coordination technique aiming at reducing the bullwhip effect [10, 11]. This technique is based on two principles:

- Lot-for-lot orders eliminate the bullwhip effect. Lot-for-lot ordering means that each company orders what is demanded of it: if its client wants 10 products, the company places an order for 10 products. With such a strategy, the bullwhip effect is eliminated, but inventory levels are not managed. Therefore, we kept lot-for-lot orders, but we added another piece of information to manage inventory levels. Precisely, orders are now vectors (O, Θ) instead of a number X, such as $X = O + \Theta$, and where O and Θ are hidden. O in (O, Θ) is set according to the lot-for-lot scheme. Therefore, the retailer transmits the market consumption to the wholesaler in O, next the wholesaler transmits this information to the pulp mill, ..., and therefore the sharing of demand information is achieved.
- Companies should react only once to each market consumption change. This second principle determines the method for choosing Θ in (O, Θ): Θ is equal to zero all the time, except when market consumption changes, in which case companies react to this change by sending non-zero Θ in order to stabilize their inventory to the initial level. Θ may be negative.

In our previous work [10, 11], we have compared two versions of our ordering scheme. In the first version, companies only have O to know the market consumption. Therefore Θ is proportional to the variation of O in respect the second principle. This first version was called "Experiment B" in [10, 11]; we now call it "ordering scheme β" (cf. Figure 2). In the second version of our ordering scheme, information centralization is used, that is, retailers multicast the market consumption to the whole supply chain. Information sharing with information centralization is much

Ordering scheme	Order O placed ($O\,p$)	Order Θ placed ($\Theta\,p$)
α	Parameter S - Inventory level (When (Inventory level $<$ Parameter s) becomes true)	0
β	Incoming order	Incoming Θ + $\lambda * \Delta$ Incoming order
γ	Market consumption	Incoming Θ + $\lambda * \Delta$ Market consumption

Figure 2. The three ordering schemes.

quicker than information sharing with (O, Θ) orders, because the market consumption transmitted in O is as slow as orders, while information centralization is assumed to be instantaneous. The second version of our ordering scheme is made more efficient by setting Θ proportional to the variation of the market consumption: as soon as the market consumption changes, non-zero Θ are sent by all companies. Moreover, companies also base O on the market consumption transmitted by retailers instead of on incoming O, again in order to react quicker to the market consumption change. This second version was called "Experiment D" in [10, 11]; we refer now to it as "ordering scheme γ" (cf. Figure 2). Finally, "ordering scheme α" is a benchmark to compare the efficiency of information sharing in the two other ordering schemes. That is, Scheme α does not use information sharing. In fact, it is a (s, S) ordering policy, which is a classic rule in Inventory Management where a company orders $(S-I)$ items when its inventory level I falls below s. These three schemes are presented in formal notations in Subsection 4.2.

2.3. Incentives for Collaboration

The two versions of our solution to the bullwhip effect, as previously presented, assume that companies share information. We now want to know whether all companies in a forest supply chain *have an incentive to use this solution*. As such a solution requires information sharing, we also want to know if all companies have an incentive to collaborate, or conversely, whether companies that would like the whole supply chain to collaborate have to pay the other companies to provide an incentive to collaborate, e.g., the companies preferring collaboration could buy the information required to collaborate. In fact, each company could prefer that the entire supply chain, except itself, collaborates. To address this issue, we consider three levels of collaboration. Each level of collaboration is represented by an ordering rule: Scheme α requires no collaboration, i.e., no information sharing, Scheme β requires basic information sharing when companies use (O, Θ) orders, and Scheme γ requires improved information sharing when companies use both (O, Θ) orders and information centralization. Figure 2 exhibits these three ordering schemes.

3. Supply Chain Management and Game Theory

Supply Chain Management can be defined as a "set of approaches utilized to efficiently integrate suppliers, manufacturers, warehouses, and stores, so that merchandise is produced and distributed at the right quantities, to the right locations, and at the right time, in order to minimize system wide costs while satisfying service level requirements" [14]. To achieve this, some researchers have used Game Theory to study incentive in supply chain management. In particular, Cachon and Netessine [2] give an overview of such studies, and note the "recent explosion of game-theoretic papers in Supply Chain Management". Our work belongs to this new area, except that we replace an analytical model by a multi-agent model that we simulate. In the same way, some authors from the Multi-Agent System field (e.g., [1, 12, 13]) suggest the use of Game Theory to analyze the behavior of a multi-agent system.

The simulation outcomes of our multi-agent simulation are used to build a game in the normal form [2] that consists of (i) players, companies or agents represented by $i = 1, \ldots, 6$, who are here the six companies, (ii) strategies r^i, which are the three ordering rules α, β, and γ available to each company i (we interchangeably use the terms: rule, scheme and strategy), (iii) payoffs/utilities, which are here replaced by company cost C^i. In fact, we do not consider a company-agent's utility, but its inventory holding and backorder costs. As we assume production cost is equal to zero, we could refer to profit by subtracting zero from the inventory and backorder cost, but instead of that, we only consider costs in order to remove negative values. Therefore, agents do not want to maximize their utility, but instead, they seek to minimize their costs. The main adaptations between our notations and traditional Economic definitions come from this difference, and other adaptations come from the use of Supply Chain Management vocabulary. We now introduce some notations used throughout the paper, then some solution concepts that are adapted from Game Theory. In these notations, we call "common parameters" the parameters shared by several companies:

$Oi^i_w =$ company i's Incoming Order O in week w;
$Op^i_w =$ company i's Placed Order O in week w;
$\Theta i^i_w =$ company i's Incoming order Θ in week w;
$\Theta p^i_w =$ company i's Placed order Θ in week w;
$I^i_w =$ company i's Inventory (or backorder when negative) in week w;
r$^i =$ strategy/ordering scheme/ordering $r$ule used by company i ;
$s =$ common parameter of the ordering scheme α;
$S =$ common parameter of the ordering scheme α;
$\lambda =$ common parameter that rules the emission of Θ in ordering schemes β and γ;
$C^i =$ company i's Cost in a fifty week simulation;
$C =$ supply chain Cost in a fifty week simulation.

In our convention, companies are written as power, e.g., C^i is company i's cost (correspondence between i and the position in the supply chain is given in Figure 3, e.g., $i = 4$ corresponds to the Paper Wholesaler). Next, we calculate company i's cost C^i as the sum of its inventory (inventory cost is $\$1.week^{-1}.unit^{-1}$) plus two times the sum of its backorder (backorder cost is $\$2.week^{-1}.unit^{-1}$) during the whole simulation (Equation 1 below). The cost of the entire supply chain C is the sum of all C^i (Equation 2 below). Next, something to the power $(-i)$ means this something for everyone except i. For example, C^{-i} is the cost of all companies except i (Equation 3 below). In other words, the overall cost of a supply chain C is equal to the sum $C^i + C^{-i}$ for any company i.

$$\forall i, C^i = \sum_{w=1}^{50} \{Max(0, I^i_w) + 2 * Max(0, -I^i_w)\} \qquad (1)$$

$$C = \sum_{i=1}^{6} C^i \qquad (2)$$

$$C^{-i} = \sum_{k \neq i} C^k \qquad (3)$$

We now adapt concepts from Game Theory to our needs (based on Jehle and Reny [5]):

Joint strategy/rule : The set r of ordering schemes $(r^1, r^2, r^3, r^4, r^5, r^6)$ used by companies is a joint strategy. Precisely, r is a vector of six strategies/rules, where the first rule refers to the Lumber Retailer's rule, the second to the Paper retailer's, ..., and the sixth to the Sawmill's rule. We can now expand the previous Equation 2 as $C(r) = \sum_{i=1}^{6} C^i(r^i, r^{-i})$, where, for example, $C^4(r^4, r^{-4}) = C^4(r^4, (r^1, r^2, r^3, r^5, r^6))$ is the cost for the paper wholesaler $(i = 4)$ of ordering with r^4 when the rest of the supply chain uses the joint strategy/rule $r^{-4} = (r^1, r^2, r^3, r^5, r^6)$. Finally, we write $r = (r^i, r^{-i})$ for any company i.

Strictly dominant strategy/rule : A strategy/rule $\hat{r}^i$ for company i is strictly dominant if $C^i(\hat{r}^i, r^{-i}) < C^i(r^i, r^{-i})$ for all (r^i, r^{-i}) such as $r^i \neq \hat{r}^i$.

In other words, a strictly dominant strategy $\hat{r}^i$ always incurs a lower cost C^i for company i than any other ordering strategy/rule, whatever strategy is used by the five other companies. Precisely, the dominant strategy $\hat{r}^i$ is always the best choice for company i, and therefore this company should use it. Such best strategies are very rare, therefore we look rather for dominated strategies in our experiments.

Strictly dominated strategy : Company i's strategy $\hat{r}^i$ strictly dominates another of its strategies $\check{r}^i$, if $C^i(\hat{r}^i, r^{-i}) < C^i(\check{r}^i, r^{-i})$ for all r^{-i}. In this case, we also say that $\check{r}^i$ is strictly dominated.

This means that whatever joint rule r^{-i} is chosen by the five other companies, a dominated strategy $\check{r}^i$ incurs higher costs C^i than another rule $\hat{r}^i$. Precisely, the dominated strategy $\check{r}^i$ is always the worst one for company i, even if it is dominated by different strategies depending on what is chosen by the other companies. For example, α is dominated by β when another player chooses α, and α is dominated by γ the rest of the time: $\check{r}^i = \alpha$ is always dominated, but not always by the same $\hat{r}^i$.

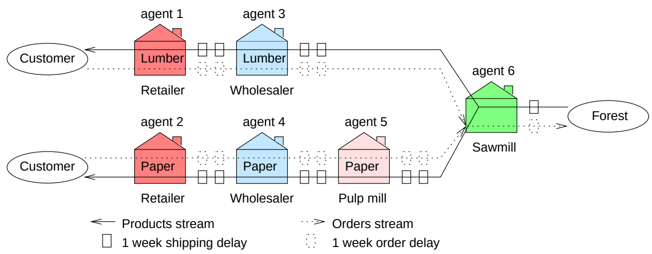

Figure 3. Model of forest supply chain used in simulations.

Pure Nash equilibrium : The joint strategy/rule $\hat{r}$ is a pure Nash equilibrium if for each company i, $C^i(\hat{r}^i, \hat{r}^{-i}) \leq C^i(r^i, \hat{r}^{-i})$ for all $r^i \neq \hat{r}^i$, where $\hat{r} = (\hat{r}^i, \hat{r}^{-n})$.

Conversely to the two above domination relations, a Nash equilibrium deals with joint strategies and not with individual strategies. A Nash equilibrium is a stable state of the supply chain: when companies choose the joint strategy $\hat{r} = (\hat{r}^1, \hat{r}^2, \hat{r}^3, \hat{r}^4, \hat{r}^5, \hat{r}^6)$ and where none of these companies has an incentive to use another ordering scheme while fully aware of the others' behavior. In other words, no company "wants to unilaterally deviate from it since such behavior would lead to" higher costs [2]. For example, if $\hat{r}^1$ is γ in the Nash equilibrium, the lumber retailer has no incentive to change for α or β when this retailer observes others' behavior. This does not mean a Nash equilibrium is the best joint strategy for the supply chain (i.e., it does not incur the minimum of C), it only means that the supply chain will remain in this equilibrium after it is reached: it is a "shaft state". Finally, to simplify the research of Nash equilibria, we first successively eliminate strictly dominated strategies, since no Nash equilibria are lost by such eliminations [8]. Next, we look for Nash equilibria in the reduced game.

Minimum of overall cost : As we consider costs instead of profits, overall cost C replaces the "social welfare" concept. As previously stated in Equation 2, the overall cost is measured as $C = \Sigma_i C^i$.

The lower the overall cost C is, the more efficient the supply chain is. If the supply chain was only one company, the goal would be to minimize C. The problem is that some companies may have no incentive to reach the minimum of C, i.e., some companies may have to sacrifice themselves by increasing their cost in order to improve the global welfare of the supply chain. Therefore, the fact that a joint strategy incurs the minimum of C means this joint strategy is the best one for the supply chain as a whole, but it does not mean that it will be used, because some companies may have an incentive to deviate from it.

Best strategy/rule : We call "best joint strategy" (it is not a Game Theory term) a joint strategy that both minimizes C and is a Nash equilibrium.

4. Simulations

The previous Game Theory concepts are used to analyze games built from simulations. We now explain how these simulations are carried out. We first introduce the simulation model based on the Québec Wood Supply Game, next we detail how the three ordering rules α, β and γ are modeled in this model, and finally the results of the analysis of simulation outcomes are given.

4.1. The Québec Wood Supply Game

Figure 3 shows how six players (human or software agents) play the Québec Wood Supply Game. The game is played by turns: each turn represents a week in reality and is played in 5 steps; these 5 steps are played in parallel by each player. In the first step, players receive their inventory (these products were sent two weeks earlier by their supplier, because there is a two-week shipping delay) and advance shipping delays between suppliers and their customers. Then in the second step, players look at their incoming orders and try to fill them. If they have backorders, they try to fill those as well. If they do not have enough inventory, they ship as much as they can, and add the rest to their backorders. In the third step, players record their inventory or backorders. In the fourth step, players advance the order slips. In the last step, players place an order to their supplier(s) and record this order. To decide what to place as orders, players compare their incoming orders with their inventory/backorder level (in our experiments, company-agents only use the ordering scheme α, β or γ). The decision that would reduce the bullwhip effect has to be taken here. Finally, a new week begins with a new step 1, and so on. Each position is played in the same way, except the Sawmill-agent: this agent receives two orders (one from the Lumber Wholesaler, another from the Pulp mill) that have to be aggregated when placing an order to the forest. The Sawmill-agent can evaluate its order by basing it on the lumber demand or on the paper demand: in our experiments, the Sawmill places an order equal to the mean of these two possible orders. Moreover, the Sawmill receives one type of product and each unit of this product generates two units: a lumber and a paper unit. That is, each incoming unit is cut in two: one piece goes to the Sawmill's lumber inventory, the other goes to its paper inventory. The full formal description of the simulation is detailed in [11].

4.2. The Three Ordering Schemes

The ordering rule r^i used by the company i in this simulation is such as $r^i \in \{\alpha, \beta, \gamma\}$, where ordering scheme α, β and γ are as follows:

- *Scheme* α: there is no collaboration here, because each company places orders on its own, neither taking into account the rest of the supply, nor sharing any information. In our experiments, we used an (s, S) ordering policy, that is, when inventory level I falls below s, the company orders $S - I$ items. We adapt in [9] the optimization model for the (s, S) ordering rule to our simulation. We find that using $s = 0$ and $S = Oi^i_w$ is optimal. As we will see with experimental results, collaboration-based Schemes β and γ perform better than the optimized α. In fact, Rule α, like the (s, S) ordering policy, is optimized under the assumption that the company using this policy has an infinite source of products as supplier. The problem is that stockouts may arise by this supplier during simulations. Therefore, the mathematical model on which the α and (s, S) ordering rules are based should be adapted to take into account an entire supply chain, rather than a single company, but this is a very difficult task. Therefore, we take $(s, S) = (0, Oi^{\,i}_{\,w})$ to calculate Op^i_w (Equation 4):

$$Op^i_w = \begin{cases} S - I^i_w & \text{if } I^i_w < s \\ 0 & \text{else} \end{cases} \quad (4)$$

  As the company applying Rule α does not use Θ, placed Θ are always equal to zero (Equation 5):

$$\theta p^i_w = 0 \quad (5)$$

- *Scheme* β: a first step of collaboration is achieved here with market consumption sharing between each company and its supplier. To share such information, companies place two-dimension orders (O, Θ) where O is the market consumption transmitted from company to company and Θ is chosen such as $O + \Theta$ represents what the company needs and such as $\Theta = 0$ when O is steady. Scheme β was proposed in [10] to reduce the bullwhip effect by stabilizing the order stream as much as possible.

  Company i's material requirements are added to incoming Θ and this sum is sent as Θ to the supplier (Equation 6 below). We have taken: $\lambda * (Oi^{\,i}_{\,w-1} - Oi^{\,i}_{\,w})$

as an estimation of material requirements. It corresponds to the inventory decrease caused by the variation of incoming order O (Oi), that is, caused by the market consumption variation. As in [10, 11], we set $\lambda = 4$ in order to have steady inventory equal to initial inventory level when market consumption becomes constant.

$$\theta p_w^i = \theta i_w^i - \lambda * (Oi_{w-1}^i - Oi_w^i) \quad (6)$$

Incoming O, which is the market consumption when all clients use Scheme β or Scheme γ, is transmitted to the client (Equation 7 below):

$$Op_w^i = Oi_w^i \quad (7)$$

- *Scheme* γ: the highest level of collaboration is achieved here. We assume there is information centralization, that is, each firm in the supply chain is provided with complete information on the actual customer consumption. Again, (O, Θ) orders are used, but now O and Θ are managed differently, but still according to the two principles outlined in Subsection 2.2.

  Scheme γ is very similar to Scheme β and was also proposed in [10]. A practical problem in the simulation is to adapt the three ordering schemes α, β and γ to make them work together. In particular, the only difference between β and γ is in the way to place O and Θ: they are based on market consumption when information centralization is achieved by the retailer, else they are based on the wholesaler's incoming order O when the retailer does not multicast its incoming O but the wholesaler does, . . ., else the company uses its own incoming order when none of its clients multicast their incoming orders. Equation 8 illustrates the case of the Pulp mill (note that Oi_w^2 is the market consumption). In the same way, Θp_w^5 is derived from Scheme β. For example, when we take again the case of the Pulp mill, the value of i in Oi_w^i and Oi_{w-1}^i in the previous Equation 6 has to be set either to 2, 4 or 5, depending on whether the paper retailer or the wholesaler broadcast or not their incoming orders O (Oi). Here, $\lambda = 2$ to have steady inventories equal to their initial levels.

$$Op_w^5 = \begin{cases} Oi_w^2 & \text{if the paper retailer multicasts its incoming orders;} \\ Oi_w^4 & \text{if the paper retailer does not multicast its incoming orders, but the paper wholesaler does.} \\ Oi_w^5 & \text{if no company multicast its incoming orders.} \end{cases} \quad (8)$$

Once these three ordering schemes are made compatible, we run our model in a multi-agent simulation where each company is an agent that uses one of these three ordering rules. We run the simulation $3^6 = 729$ times, with $I_0^i = 0$ as initial condition for every agent i. For each run, we use another combination of the three ordering rules among the six agents: all companies use Scheme α in the first run, then all companies use Scheme α except one company which uses Scheme β in the second run, . . ., and finally, all companies use Scheme γ in the 729$^{\text{th}}$ simulation. For each run of this simulation, each company-agent inventory holding and backorder cost is evaluated. In this paper, we do not focus on the reduction of the bullwhip effect, which is measured as the standard-deviation σ of companies' orders and not as a cost, but on the *financial benefits of reducing the bullwhip effect* in a collaborative way. The outcomes (i.e., each company's cost) of the 729 simulations are then used to build a game in the normal form. We next analyze this game to find dominated ordering schemes, Nash equilibria and the minimum of C.

4.3. Results

The analyze of this huge game (6-dimension matrix filled with 6 ∗ 729 individual costs) is carried out with Gambit 0.97.0.4 [8], a software distributed under the GNU General Public License. First, α is dominated for Paper and Lumber Retailers. Next, Figures 4 and 5 show the results of this analysis. The format of data in Figure 4 is $(r^1, r^2, r^3, r^4, r^5, r^6) \rightarrow C^1 + C^2 + C^3 + C^4 + C^5 + C^6 = C$. For instance, the next-to-last Nash equilibrium, called JS^{eq4} is $(\gamma, \beta, \gamma, \beta, \beta, \beta) \rightarrow 1,008 + 1,740 + 888 + 1,644 + 1,356 + 1,752 = \$8,388$, in which, for example, the Pulp mill $(i = 5)$ uses $r^5 = \beta$, which incurs for this company a cost $C^5 = \$1,356$. With this notation, the two joint strategies in JS^{min} incur the mininimum of the supply chain cost $C = \$5,100$, and every company i has the same individual cost C^i in them. Furthermore, it is very interesting to note that these two joint strategies are also Nash equilibria, i.e., $JS^{min} = JS^{eq1}$. We earlier called these two joint strategies the best ones, because they are they incur the lowest supply chain cost, while no company has an incentive to deviate from them. In addition to the two Nash equilibria JS^{eq1}, there are four other equilibria $JS^{eq2}, JS^{eq3}, JS^{eq4}$ and JS^{eq5}. The ordering rule α does not appear in any of these equilibria, which shows that companies prefer collaboration. In the same way, γ is used by almost every company in $JS^{min} = JS^{eq1}$ in which the supply chain cost C is minimal, while this supply chain cost C is higher when several companies use β instead of γ (cf. JS^{eq4} and $JS^{eq5} = JS^{hom\beta}$). This makes us think that C should be minimised when all companies use γ. In fact, full collaboration should be better for the supply chain than basic collaboration. Surprisingly, this is not the case, because a homogeneous supply chain using γ incurs $C = \$6,552$ (cf. $JS^{hom\gamma}$), while a heterogeneous supply chain using JS^{eq1} or JS^{eq2} incurs a lower C. Here, we call "heterogeneous" a supply chain in which every company has the same behaviour, i.e., every company uses the same ordering scheme. Finally, with a homogeneous supply chain us-

	$(\mathbf{r^1}, \mathbf{r^2}, \mathbf{r^3}, \mathbf{r^4}, \mathbf{r^5}, \mathbf{r^6}) \to \mathbf{C^1} + \mathbf{C^2} + \mathbf{C^3} + \mathbf{C^4} + \mathbf{C^5} + \mathbf{C^6} = \mathbf{C}$
Minimum of C	$JS^{min} = \{(\gamma,\gamma,\beta,\gamma,\gamma,\gamma),(\gamma,\gamma,\gamma,\gamma,\beta,\gamma)\} \to 696 + 1,068 + 576 + 948 + 720 + 1,092 = 5,100\$$
Nash equilibria	$JS^{eq1} = \{(\gamma,\gamma,\beta,\gamma,\gamma,\gamma),(\gamma,\gamma,\gamma,\gamma,\beta,\gamma)\} \to 696 + 1,068 + 576 + 948 + 720 + 1,092 = 5,100\$$
	$JS^{eq2} = (\gamma,\gamma,\gamma,\beta,\gamma,\gamma) \to 720 + 1,092 + 600 + 972 + 660 + 2,136 = 6,180\$$
	$JS^{eq3} = (\beta,\gamma,\beta,\beta,\gamma,\beta) \to 894 + 1,266 + 774 + 1,146 + 834 + 2,400 = 7,314\$$
	$JS^{eq4} = (\gamma,\beta,\gamma,\beta,\beta,\beta) \to 1,008 + 1,740 + 888 + 1,644 + 1,356 + 1,752 = 8,388\$$
	$JS^{eq5} = (\beta,\beta,\beta,\beta,\beta,\beta) \to 1,098 + 1,830 + 978 + 1,734 + 1,446 + 1,848 = 8,934\$$
Homogeneous supply chain	$JS^{hom\alpha} = (\alpha,\alpha,\alpha,\alpha,\alpha,\alpha) \to 9,682 + 15,204 + 54,186 + 90,998 + 254,294 + 724,130 = 1,148,494\$$
	$JS^{hom\beta} = (\beta,\beta,\beta,\beta,\beta,\beta) \to 1,098 + 1,830 + 978 + 1,734 + 1,446 + 1,848 = 8,934\$$
	$JS^{hom\gamma} = (\gamma,\gamma,\gamma,\gamma,\gamma,\gamma) \to 696 + 1,428 + 576 + 1,332 + 1,128 + 1,392 = 6,552\$$

Figure 4. Experimental results (raw data).

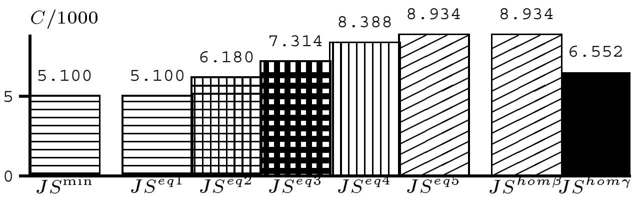

Figure 5. Experimental results (histograms).

ing $JS^{hom\alpha}$, i.e., when every company uses α, all individual costs C^i and the supply chain cost C are much higher than with β and γ. This shows that optimizing for single companies can be outperformed by considering the whole supply chain in order to propose to companies some ways of collaboration, i.e., information sharing in our case. This very high cost C is the reason why $JS^{hom\alpha}$ is not mentionned in Figure 5

5. Discussion

First, we only consider in this paper pure strategies, that is, companies use the same ordering rule in each of the fifty weeks of a simulation. The concept of mixed strategies also exists in Game Theory, in which there is the addition of probability on the usage of the rules (e.g. the paper retailer uses α 75% of the time and β the remaining time). In fact, mixed strategies cannot be used in our simulations, because of the two following remarks:

1. Conversely to traditional Game Theory, we cannot determine the expected outcome of a simulation with mixed strategies based on two simulations with pure strategies. The main reason is that there is a transition period in the supply chain when a company switches from one ordering rule to another. For example, when the paper retailer switches from α to β, products previously ordered with α will still arrive after four weeks, which has an impact on the paper retailer's inventory level and behavior. This problem of "state transition" does not exist in traditional Game Theory. If we want to approximate the expected outcome for some specific mixed strategy, we have to simulate this mixed strategy over a very large number of weeks.
2. We do not know of any algorithms that determine Nash equilibria in mixed strategies in a reasonable time for games as large as ours. The determination of Nash equilibria in pure strategies requires the comparison of individual outcomes, while mixed strategies requires the resolution of linear equations, a more complex task.

Here, we are concerned with the interpretation of the analysis of the simulation outcomes. These results are not obtained with a dynamic simulation, that is, a simulation in which companies are allowed to change their behavior. On the contrary, each company keeps the same ordering rule during the fifty weeks of the run of the simulation. On the one hand, if companies were able to change of ordering rule during a run, i.e., if we changed the simulation for a dynamic one, the supply chain would either stabilize on one of the Nash equilibria that we found, or will never stabilize. This possible stabilization would depend (i) on the initial state of the supply chain (initial inventory levels, orders and shippings) and (ii) on the decision process used by companies to change their ordering scheme. On the other hand, we can propose another interpretation of our simulations in which companies negotiate before playing. This assumption does not totally conform with Game Theory, because players only have one choice to make, but it corresponds with our model and with real life. Here, equilibria correspond to the outcome of the negotiation made before the first week of simulation. In other words, before producing any products, the six companies must first agree on which joint strategy they want to reach. In this negotiation, $JS^{eq1} = JS^{min}$ can hold, because the negotiation can only conclude on a Nash equilibrium. Companies should try to reach this equilibrium instead of another one. After the negotiation has ended, i.e., after a Nash equilibrium has been chosen, companies have to sign contracts allowing them to use the rules β and γ, because information shared in rules β and γ have to be kept secret, and both information sharing and centralization requires the installation of some supporting technologies, in particular, based on the Internet. By definition, no other contract forcing companies not to deviate from the chosen equilibrium is required.

Finally, we insist on the interpretation of outcomes in our experiments, and in particular our Nash equilibria. Such outcomes represent costs for fifty week simulations. In other words, when companies consider them, they only base their decision on long term costs. This means that Nash equilibria are stable states for the supply chain if all companies use the same time horizon in their decision process; if a company considers a shorter horizon, it could have an incentive to leave such equilibria. On the other hand, collaboration is often considered on a long horizon, because it requires companies to sign contracts for secrecy agreements and to install some collaboration-support devices: our time horizon is therefore well chosen.

6. Conclusion

This paper studied companies' incentive to collaborate in order to reduce the bullwhip effect. Collaboration is seen here as information sharing. Precisely, we have designed an ordering scheme [10, 11] to reduce the bullwhip effect, a phenomenon in which demand variability amplifies in the supply chain and which causes costs due to higher inventory levels and agility reduction. We compared two versions of this ordering scheme with the traditional (s, S) ordering policy, usually used in Inventory Management. That is, we considered three levels of collaboration: (α) no collaboration when companies use a classic (s, S) ordering policy, (β) little collaboration when companies use our basic ordering scheme, which is based on information sharing between each company and its suppliers, and (γ) full collaboration when companies enhance our ordering scheme with information centralization, that is, companies receive the actual market consumption from retailers in real-time.

Three main conclusions have been drawn from our experimental results. First, the two joint strategies that minimize the supply chain cost are Nash equilibria. Therefore, no company has an incentive to deviate from them. In these two equilibria, almost every company fully collaborates, which recalls that collaborating, through information sharing, or better, through the centralization of information, is always seen as a good practice in a supply chain. Second, several Nash equilibria exist in the considered supply chain and all of them require (basic or full) collaboration. Third, there are no equilibria where one or several companies do not collaborate.

As future work, we will first change the market consumption pattern in order to verify if the results still hold. Next, we will look for relations of Pareto-dominance in simulation outcomes to order the Nash equilibria. Finally, we will adapt our model to assume companies want to maximize their utility/profit, instead of minimizing their cost. Therefore, costs will take into account production and inventory activities. The supply chain will earn money each time a retailer sells a product to the market, and loose sales when market consumption exceeds the quantity of products available by retailers. Backorders in the supply will no longer cost money: the only goal of each company will be, for retailers, to have enough products. Therefore, companies' goals will be quite different: instead of aiming at zero inventory and zero backorder, companies will try to maximize the retailer's sales, while minimizing their own inventory.

7. Acknowledgment

We would like to thank FOR@C, the Research Consortium in E-business in the forest products industry (Université Laval, Québec, Canada), for supporting this research. This work was also partially supported by the National Sciences and Engineering Research Council of Canada (NSERC).

References

[1] C. Boutilier, Y. Shoham, and M. P. Wellman. Economic principles on multi-agent systems (Editorial). *Artificial Intelligence*, 94(1-2), 1994.

[2] G. P. Cachon and S. Netessine. Game Theory in Supply Chain Analysis. In D. Simchi-Levi, S. D. Wu, and Z.-J. M. Shen, editors, *Supply Chain Analysis in the eBusiness Era*. Kluwer (to appear), 2003.

[3] C. Carlsson and R. Fullér. Reducing the bullwhip effect by means of intelligent, soft computing methods. In *Proc.34th Hawaii International Conference on System Sciences*, 2001.

[4] F. Chen, Z. Drezner, J. K. Ryan, and D. Simchi-Levi. Quantifying the bullwhip effect in a simple supply chain: The impact of forecasting, lead times, and information. *Management Science*, 46(3):436–443, 2000.

[5] G. A. Jehle and P. J. Reny. *Advanced Microeconomic Theory*. Addison Wesley, 2000. 2d Edition.

[6] H. L. Lee, V. Padmanabhan, and S. Whang. The bullwhip effect in supply chain. *Sloan Management Review*, pages 93–102, 1997.

[7] H. L. Lee, V. Padmanabhan, and S. Whang. Information distortion in a supply chain: the bullwhip effect. *Management Science*, 43(4):546–558, 1997.

[8] R. D. McKelvey, A. M. McLennan, and T. L. Turocy. *Gambit 0.97.0.3 manual*. 2003. http://econweb.tamu.edu/gambit/manual/gambit-manual.pdf (accessed 23 October 2003).

[9] T. Moyaux. *Multi-Agent Simulation and Analysis of Collaborative Strategies in the Québec Forest Supply Chain*. Ph.D thesis, Univ. Laval (Québec City, Québec, Canada), 2004.

[10] T. Moyaux, B. Chaib-draa, and S. D'Amours. Multi-agent coordination based on tokens: Reduction of the bullwhip effect in a forest supply chain. In *Proc. 2nd int. joint conf. on Autonomous Agents and MultiAgent Systems (AAMAS'03)*, Melbourne, Australia, 2003.

[11] T. Moyaux, B. Chaib-draa, and S. D'Amours. The impact of information sharing on the efficiency of an ordering approach in reducing the bullwhip effect. *IEEE Trans. on Systems, Man, and Cybernetics*, 2004 (submitted).

[12] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers*. MIT Press, Cambridge, Massachusetts, 1994.

[13] T. W. Sandholm. Distributed rational decision making. In G. Weiss, editor, *Multiagent Systems - A Modern Approach to Distributed Artificial Intelligence*, pages 201–258. MIT Press, Cambridge, Massachusetts, 1999.

[14] D. Simchi-Levi, P. Kaminsky, and E. Simchi-Levi. *Designing and managing the supply chain*. McGraw-Hill Higher Education, 2000.