Information Sharing as a Coordination Mechanism for Reducing the Bullwhip Effect in a Supply Chain

Thierry Moyaux, Brahim Chaib-draa, *Senior Member, IEEE*, and Sophie D'Amours

***Abstract*—The bullwhip effect is an amplification of the variability of the orders placed by companies in a supply chain. This variability reduces the efficiency of supply chains, since it incurs costs due to higher inventory levels and supply chain agility reduction. Eliminating the bullwhip effect is surely simple; every company just has to order following the market demand, i.e., each company should use a lot-for-lot type of ordering policy. However, many reasons, such as inventory management, lot-sizing, and market, supply, or operation uncertainties, motivate companies not to use this strategy. Therefore, the bullwhip effect cannot be totally eliminated. However, it can be reduced by information sharing, which is the form of collaboration considered in this paper. More precisely, we study how to separate demand into original demand and adjustments. We describe two principles explaining how to use the shared information to reduce the amplification of order variability induced by lead times, which we propose as a cause of the effect. Simulations confirm the value of these two principles with regard to costs and customer service levels.**

***Index Terms*—Agents, bullwhip effect, coordination mechanisms, multiagent systems, supply chain management.**

I. Introduction

Supply chain management can be defined as a set of techniques utilized to efficiently integrate suppliers, manufacturers, warehouses, and stores, so that merchandize is produced and distributed in the right quantities, to the right locations, and at the right time, in order to minimize system wide costs while satisfying the service-level requirements [1].

As we can see, supply chains are distributed systems, and thus, issues of stream fluctuations may appear therein. In fact, in the case of supply chains, this issue of fluctuations is known as the bullwhip effect (or Forrester's effect [2]). This effect is a problem of coordination consisting of an amplification of demand variability in the supply chain, so that raw material producers receive orders that are more variable and unpredictable than that of the retailers. Basically, this problem leads to unnecessary inventory and decreased customer service levels due to backorders, i.e., to inventory shortages/lost sales. We present the bullwhip effect and review the relevant literature in Section II.

Manuscript received October 13, 2004; revised June 20, 2005, and September 29, 2005. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada. This paper was recommended by Associate Editor J. Chen.

T. Moyaux was with the Department of Computer Science and Software Engineering, Université Laval, Quebec City, QC, G1K 7P4, Canada. He is now with the Department of Computer Science, University of Liverpool, Liverpool L69 3BX, U.K. (e-mail: moyaux@csc.liv.ac.uk).

B. Chaib-draa is with the Department of Computer Science and Software Engineering, Université Laval, Quebec City, QC, G1K 7P4, Canada (e-mail: chaib@ift.ulaval.ca).

S. D. Amours is with the Department of Mechanical Engineering, Université Laval, Quebec City, QC, G1K 7P4, Canada (e-mail: sophie.damours@forac.ulaval.ca).

Digital Object Identifier 10.1109/TSMCC.2006.887014

We next study how the supply chain should behave to reduce the bullwhip effect. Our ultimate goal is to improve supply chain efficiency by reducing this effect of amplification of demand variability, while keeping a low inventory and adequate customer-service levels. Since we see each company as a reactive agent that applies a given ordering rule, our goal is to stabilize flows in this multiagent system, i.e., to have a stable supply chain. In this study, we consider a split supply chain where companies have only one supplier and deal with one product. With this model, we show why lead times induce the bullwhip effect, although previously lead times have only been seen as an aggravating factor of another cause of this effect. The solution that we propose to this specific cause is based on demand information sharing, which is often suggested to reduce the bullwhip effect [1], [3]–[5]. Lot-for-lot ordering (that is, each company orders what is demanded by its clients) is known to eliminate the bullwhip effect, and propagates the market consumption in the supply chain. Unfortunately, this often results in backorders due to inventory variations. This is why we propose to distinguish between real market demand from demand variation required to stabilize inventory levels, i.e., between the original demand and the required adjustments. To achieve that, every company places vectors of orders $[O, \Theta]$, where O's are the market consumption transmitted from company to company and Θ's represent the difference with O's with the needs of the companies. The idea is that all companies would order O's if there were no bullwhip effect, but they do not order O for many reasons, in particular, inventory variation. The information about market consumption O is used by companies in the supply chain as the reference point from which their orders must deviate as little as possible: thus, Θ's "encapsulate" the bullwhip effect. Section III explains why and how to separate the original demand from adjustments so that the bullwhip effect is reduced, in particular, through the proposition of two principles ruling O and Θ.

This solution is compared experimentally to some other ordering schemes by using the Québec wood supply game (QWSG). This game is derived from the Wood Supply Game [6], [7] to Québec (a Canadian province) wood industry specificities. These games are adaptations of the beer game [8] for the wood industry, and were designed to make players aware of supply chain dynamics, in particular, the bullwhip effect. We implemented the QWSG as a software multiagent system [9] and simulated it with 19 market consumption patterns and under seven different ordering schemes. Three of these schemes use 2-D $[O, \Theta]$ orders, while the four others use classic reorder schemes with one dimension $X(= O + \Theta)$. The QWSG, its multiagent implementation, experimental results, and a discussion are provided in Section IV. An important point of the discussion is that the proposed ordering scheme inverses company interest. In fact, conversely to the classic bullwhip

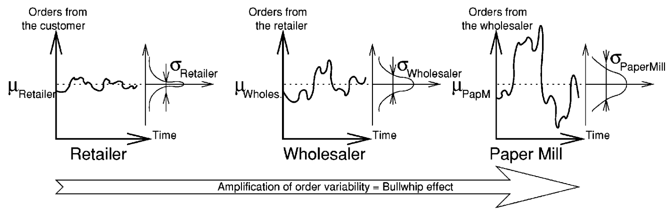

Fig. 1. Bullwhip effect [4], [5].

TABLE I
PROPOSED CAUSES AND SOLUTIONS OF THE BULLWHIP EFFECT

Causes	**Proposed solutions**	**Authors**
Demand forecast updating	Information sharing (e.g., VMI, CRP...), echelon-based inventory and leadtime reduction	[4], [5]
Order batching	EDI (Electronic Data Interchange) and Internet technologies	[4], [5] [11]
Price fluctuation	EDLP (Every Day Low Pricing)	[4], [5]
Rationing and shortage gaming	Allocation based on past sales	[4], [5]
Misperception of feedbacks	Giving a better understanding of the supply chain dynamics to managers	[12], [13] [2], [14] [8]
Local optimization without global vision	None	[3], [15] [16] [1]
Company processes	None	[11]

effect, which amplifies along the supply chain (most upstream suppliers are more disturbed by this effect than retailers), ordering schemes following our two principles induce inventory variations that last longer at retail sites than at upstream sites.

II. BACKGROUND AND LITERATURE REVIEW ON THE BULLWHIP EFFECT

We now present the bullwhip effect. Fig. 1 shows how this effect propagates in a simple supply chain with only three companies: a retailer, a wholesaler, and a paper mill. In this figure, the retailer exclusively sells to the customer and buys from the wholesaler, the wholesaler sells to the retailer and buys from the paper mill, and the paper mill sells to the retailer and buys from an unknown supplier. The ordering patterns of the three companies are similar in the way that the variabilities of an upstream site are always greater than those of the downstream site [5]. As a variability, the bullwhip effect is measured by the standard deviation σ of orders. Note that the means μ of orders are all equal in our example given in Fig. 1.

A. Causes Sustaining the Bullwhip Effect

There are several consequences of the bullwhip effect. In a few words, this effect incurs costs due to: 1) higher inventory levels; 2) supply chain agility reduction; 3) decrease of customer service levels; 4) ineffective transportation; 5) missed production schedules [10]; and 6) stockpiling due to a high degree of demand uncertainties and variabilities [4]. Table I summarizes the causes and solutions of the bullwhip effect proposed in the literature, which are now detailed. Lee *et al.* [4], [5] proposed the first four causes and solutions.

1) Demand Forecast Updating: Companies base their orders on forecasts, which are themselves based on their incoming orders, although such forecasts are not perfectly accurate. Therefore, companies order more or less than what they really require to fulfill their demand. In other words, forecasting errors amplify the variability of orders. A solution proposed to this cause is information sharing: each client provides more complete information to its supplier in order to allow the supplier to improve its forecasting. Information sharing is already part of industry practices, such as vendor-managed inventory (VMI), continuous replenishment program (CRP), etc.

2) Order Batching—Lot Sizing in A More General Way: Companies discretize orders for profiting from economies of scales, and therefore, place orders for more or less products than what they actually need.

The proposed solution to this cause is electronic transactions (e-commerce, EDI, etc.) to reduce transaction costs and thus make companies' orders more frequent and for smaller quantities. Similarly, the size of production batches may be reduced with single minute exchange of die (SMED), which may next reduce the quantities ordered.

3) Price Fluctuation: Every client (company or end-customer) profits from promotions by buying more products than what it really requires, and next, buying nothing when the promotion stops because it has enough products in inventory. The proposed solution is the every day low pricing (EDLP) policy, where price is set at the promotion level. However, EDLP also has some drawbacks, e.g., always looking for the lowest price may put a stress on the supply chain that may eventually reduce profits [17].

4) Rationing and Shortage Gaming: Since every client has opportunist behavior, it overorders when its supplier cannot fulfill its entire demand, e.g., in the case where the supplier has a machine breakdown. Through such behavior, this client does not hope to receive the quantity that it has ordered, but a lower quantity that matches its actual need. Since this behavior occurs when the supplier allocates shipping in proportion to the ordered amount, it is preferable to allocate the few available products in proportion with the history of past orders.

Other authors have extended Lee and his colleagues' causes to the bullwhip effect.

5) Misperception of Feedback: Sterman [8] has noted that players in the beer game place orders in a nonoptimal way because they do not understand the whole dynamics in their supply chain. For example, they do not correctly interpret their incoming orders, and in consequence, smooth their orders when they should order more, because they do not understand that market consumption has increased.

6) Local Optimization Without Global Vision: Several authors [15], [16] have noted that companies maximize their own profit without taking into account the effect of their decisions on the rest of the supply chain. In particular, some companies use an ordering scheme such as the (s, S) policy (in which the company orders for S-I products when inventory level I falls below s) that is the operationalization of this local optimization. It has been formally proven that some of these policies induce the bullwhip effect [1], [3].

7) Company Processes: Taylor *et al.* [11] propose two causes of the bullwhip effect: variability in machine reliability and output, and variability in process capability and subsequent product quality. In these two causes, which are summarized as "company processes" in Table I, production problems at each workstation are amplified from one workstation to another. This cause recalls that intracompany problems and uncertainties may affect each company's behavior, which in turn may make them change the way they place orders.

B. Literature Review on the Bullwhip Effect

This presentation of the bullwhip effect and its proposed causes is a synthesis of many works. Basically, Scholl [18] noted that many of these works belong to system dynamics modeling and agent-based modeling, which are two prominent nonlinear modeling schools. In a broader way, we extend this classification of the literature about the bullwhip effect into three broad classes.

1) Formal Studies Relative to the Bullwhip Effect: Mathematical tools from different fields have been used. First, Forrester [2], [14] introduced the bullwhip effect in the field of *system dynamics* in 1958,[1] and other studies of this effect were next carried out in this field [19], [20].

Second, *inventory management* is the main field concerned with this effect, since it is strongly related with order placement. In particular, Lee *et al.* [4] proposed a formal description of the four first causes stated above. In the same way, Simchi-Levi *et al.* [1], [3] studied the first of those four causes (demand forecast updating) further by focusing on the impact of lead times and moving average forecasting and on how the bullwhip effect is reduced with information centralization.[2] In the same vein, Kelle and Milne [21] wrote a work similar to [1] and [3] with the (s, S) ordering policy. Third, *economists* [15], [16] have also studied how local optimizations done by companies without taking into account the rest of the supply chain cause the bullwhip effect. The difference with the previous approach is the fact that optimization is explicitly taken into account in economics models, while it disappears in classic inventory management models of the bullwhip effect: the latter approach bases its ordering policies on local optimization too, but it uses these policies without reconsidering the optimization process on which these policies are based. Fourth, *traffic flow theory* was translated into supply chain management vocabulary by Daganzo [12] in order to study how ordering policies can stabilize flows in the supply chain. More precisely, Daganzo represents the history of flow at each company with curves of cumulative count, i.e., the greater a flow is, the greater the gradient of the curve representing this flow is. With this representation of flows, he gave properties that the ordering policies ruling the order streams should have in order to avoid the bullwhip effect, and shows that all currently used policies lead to the bullwhip effect.

Finally, in an approach similar to traffic flow theory, Dejonckheere *et al.* [13] have also focused on stabilizing streams in the supply chain, except that they used another formalism called *control theory* to verify the impact of exponential smoothing algorithms for forecasting in a way similar to that of Simchi-Levi *et al.* [1], [3] and Kelle and Milne's [21].

2) Empirical Studies of the Bullwhip Effect: Instead of focusing on mathematical representations of the bullwhip effect, many researchers have studied the bullwhip effect in an empirical way. For example, Lee *et al.* [5] gave a nonformal description of their paper [4], in which the first four above causes of the bullwhip effect are described. Similarly, Wilding [22] explained in general terms that uncertainty in the supply chain is generated by three interacting sources: the bullwhip effect, deterministic chaos (chaos appears when the system is deterministic, i.e., chaos appears in the system when there is a definite rule with no random terms governing the dynamics of the system), and parallel interactions (interactions between companies in the same echelon may appear: a retailer has an influence not only on its suppliers, but also on other retailers).

In a more practical way, Fransoo and Wouters [23] proposed a method for the measurement of the bullwhip effect. In the same way, the LEAn Processing (LEAP) project [11], [24] has focused on the bullwhip effect in three echelons of the automotive component supply chain in the U.K.

3) Simulation-Based Studies of the Bullwhip Effect: Besides formal models and empirical approaches, simulation is increasingly seen as an efficient tool in supply chain management [25], [26]. In the context of the bullwhip effect, Yung and Yang [27] represented each company as an *agent* (i.e., as an autonomous, reactive, and proactive software that can interact with other agents [9]) that minimizes its costs subject to some constraints. Since these agents work in parallel, the optimization of the supply chain is done concurrently. In a similar approach based on agents, Carlsson and Fullér [10] used *fuzzy logic* to estimate demand for the upcoming period. Like Chen *et al.* [3] (mentioned above as a work in inventory management) but with the *agent* paradigm, Yan [28] has studied the impact of lead-time distribution on the bullwhip effect.

As previously stated, the bullwhip effect has also been studied using the beer game, which can be defined as a simulation of a supply chain used to teach the bullwhip effect. In this simulation, some researchers [8], [29] have looked for the managers' cognitive limitations that cause this effect of demand variability. In fact, although there are mathematical tools to manage inventories, some people still use their intuition when placing orders in real supply chains. The problem lies in the fact that people have some difficulties understanding the dynamics of a supply chain, because there are complex feedback loops, time delays, and past orders to consider together. Besides that, Kimbrough *et al.* [30] have gathered software agents and the beer game by replacing human players by agents in order to find the best ordering scheme with a genetic algorithm. Next, some modifications have been proposed to the beer game: Chen and Samroengraja [31] changed some parameters, Fjeld and Haartveit [6], [7] adapted it to the North European forest industry to study how the structure of the game can result in a mismatch between supply and demand, and FOR@C's researchers at Université

[1] However, Scholl [18] says that Sterman points out that this phenomenon was described at least as early as the 1920s and 1930s in economics and management science literatures.

[2] Information centralization is a special form of information sharing in which retailers multicast their sales to the rest of the supply chain instantaneously and in realtime.

Laval (Québec, Canada),[3] adapted this latter game in order to take Québec forest industry specificities into account. This last game is called the Québec wood supply game (QWSG).

FOR@C's researchers used the multiagent paradigm to simulate this game [32], and made it more realistic in [33]. In particular, information sharing is studied as a decentralized coordination technique in [34]; technique that we will develop in this paper.

Our problem is to find a way to place orders that is the most efficient for the whole supply chain. As a solution, we have looked for an ordering scheme that *stabilizes placed orders*, while minimizing inventory levels and avoiding stockouts. Our approach can be compared to that of Kimbrough *et al.* [30] in that we use multiagent systems to find a good ordering scheme. The main difference is that Kimbrough looked for an ordering rule that minimizes inventory and backorders for a certain period of time with a genetic algorithm (therefore, the best ordering pattern may change depending on the duration of the simulations), while we focus a good ordering rule for any duration of simulation. Another slight difference is that Kimbrough *et al.*'s work uses the Sterman's beer game [8], while we use an adaptation of this game to the Québec forest industry: the QWSG.

We now present the solution we propose to reduce the bullwhip effect while minimizing inventory levels and avoiding stockouts.

III. Information Sharing to Reduce the Bullwhip Effect

To introduce our two ordering schemes, we first comment on the cause of the bullwhip effect that it specifically addresses. In fact, we propose regarding ordering and shipping lead times as a cause of the bullwhip effect, while lead times are only seen as an aggravating factor of the cause "demand forecast updating" listed in Table I. As a consequence, ordering and shipping lead times could be added to Table I, but we will see that it is only a particular case of "misperception of feedbacks." Then, we propose two principles that an ordering scheme should have to reduce the impact of this cause. Finally, we present the behavior of the supply chain under our two principles.

A. Why Delays Cause the Bullwhip Effect in the QWSG?

Several causes of the bullwhip effect have been proposed in the literature for real supply chains, but few of these causes occur in the QWSG. From our viewpoint, only two of the causes in Table I can be found in the QWSG, namely the "demand signal processing" [4], [5] and the "misperception of feedback" [8]. Since the former of these two possible causes is related to demand forecasting, it can explain the bullwhip effect in the broad version of the QWSG, because human players intuitively forecast their future incoming demand, but not in our simulation in which orders are only based on the last demand. In fact, we assume there is no forecast because our two proposed ordering schemes base the current order on the last demand, but we can also see this method as a forecast based only on the last demand. The above second cause could explain why we still had a bullwhip effect when software agents replaced human players. In fact, players' understanding of the supply chain dynamics was not used directly in our experiments, but the ordering policies that they apply could be designed so that these dynamics are taken into account. When we looked for efficient ordering policies, we found that the cause "misperception of feedback" can be detailed as "ordering and shipping lead times," as now presented in three points with Fig. 2(a).

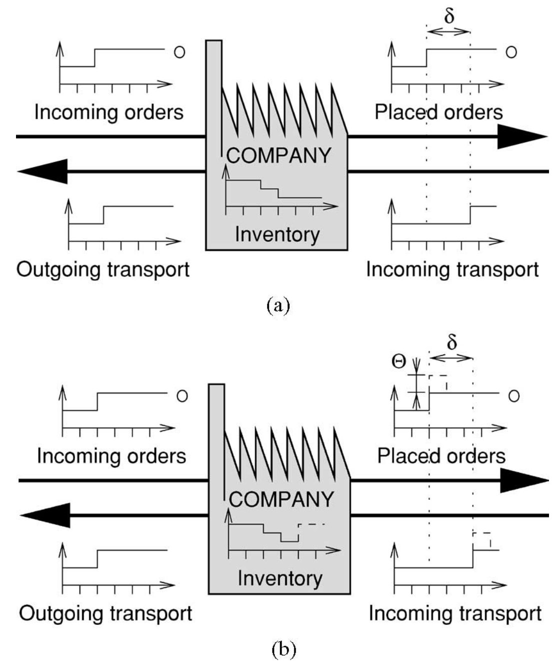

Fig. 2. Lot-for-lot ordering policy with $[O, \Theta]$ orders. (a) Lot-for-lot ordering policy. (b) $[O, \Theta]$ orders.

1) The lot-for-lot ordering policy eliminates the bullwhip effect, because each company has the same ordering pattern as its client and thus, as the market consumption. Therefore, the two curves *Incoming orders* and *Placed orders* are identical in Fig. 2(a). Since the bullwhip effect is measured as the standard deviation of placed orders, we can see that the standard deviation of each company's orders is exactly the same as the standard deviation of its client's orders, and therefore, as the standard deviation of the market consumption. This explains why a lot-for-lot ordering policy eliminates the bullwhip effect.
2) The considered company tries to fulfill its entire demand, and thus, the two curves *incoming orders* and *outgoing transport* are the same, i.e., as many products are shipped as are ordered. This is true as long as there are enough items in *inventory*, i.e., as long as no backorder occurs.
3) The curve *incoming transport* has the same pattern as the other three curves, except that it is delayed by the temporal shift δ that corresponds to the ordering and shipping lead times. The problem is that the inventory is not managed, because the temporal shift δ makes inventory decrease (respectively increase, when we inverse the pattern of *incoming orders*), due to the fact that the company ships

[3]http://www.forac.ulaval.ca

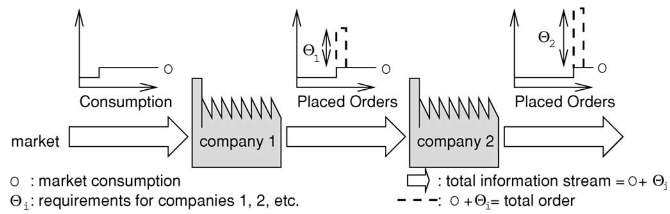

Fig. 3. Information streams cut into two parts.

more (respectively less) products than it receives during δ. Finally, notice that *incoming transport* has the same pattern as the other three curves only when the supplier has no stockouts, because the supplier is assumed to want to fulfill its entire demand, like the considered company.

Since every company wants to avoid stockouts (respectively huge inventory), *rather than eliminate* the bullwhip effect, it does not use the lot-for-lot ordering policy. Instead, it overorders (respectively underorders) in comparison with the lot-for-lot policy to stabilize its inventory, which amplifies the demand variabilities because the company overorders (respectively underorders) when the demand increases (respectively decreases).

As a result, lead times in the supply chain make it so that the bullwhip effect always appears each time the market consumption has an infinitesimal *change*, if companies want to *keep a steady inventory*. We can note here that some of the other causes of the bullwhip effect presented in Table I induce the bullwhip effect even with a steady demand, while lead times only amplify the fluctuations of orders, but do not induce fluctuations when the demand is steady.

As we can see, our problem is not only to reduce the bullwhip effect, because company agents in the QWSG have only to apply the lot-for-lot ordering policy to eliminate this effect, but we also have to manage inventories. In our solution, we propose to use the information-sharing solution presented in Fig. 2(b), in which each company uses a vector $[O, \Theta]$ of two orders (O like Orders and Θ like Tokens, as these two pieces of information were called in our previous papers, and O and Θ also have the advantage of looking similar, while they have a very similar meaning: both are Orders). Remember that O's follows the lot-for-lot policy to avoid the bullwhip effect, and Θ's are used to order more or less products than O to stabilize the inventory level. As a consequence, O transmits the market consumption information to the whole supply chain, as illustrated by Fig. 3. We now present the two principles ruling the use of O and Θ.

B. Why and How Does a Company Use $[O, \Theta]$ Orders?

Since the bullwhip effect may appear in Θ, i.e., nonzero Θ's may be emitted anytime, we now present two principles ruling O and Θ. Indeed, $[O, \Theta]$ ordering schemes should be based on the following two principles, and not only on one of them.

1) First Principle—The Lot-For-Lot Ordering Policy Eliminates the Bullwhip Effect, But Does Not Manage Inventories: This first principle indicates the manner in which O in $[O, \Theta]$ is chosen. As previously stated in Fig. 2(a), the bullwhip effect is eliminated with the lot-for-lot policy, but the problem is that inventory levels are not managed. Therefore, we keep lot-for-lot orders for ruling O, but we add another piece of information Θ to adjust inventory levels.

2) Second Principle—Companies Should React Only Once to Each Market Consumption Change: This second principle indicates the way of choosing Θ in $[O, \Theta]$. Θ's are equal to zero all the time, except when the market consumption changes, in which case companies react to this change by sending nonzero Θ's in order to stabilize their inventory to the initial level. The purpose of Θ's is to trigger a product wave from the most upstream company (e.g., the forest) when this company receives these Θ's. This product wave will increase (or decrease when $\Theta < 0$) each company's inventory as it travels the supply chain down to the retailers. The size of this wave is given by Θ's, which are the sum of positive and negative company needs, added successively by every company when it places orders. Each company transmits to its supplier the Θ's incoming from its clients and adds to them its relative requirements in comparison with the market consumption O. We now present this global behavior of the supply chain incurred by our two principles.

C. Supply Chain Behavior Under Our Two Principles

To present the global behavior that the supply chain exhibits with our two principles, we assume that there is a unique increase in market consumption, similar to Figs. 2 and 3. The goal that $[O, \Theta]$ orders must reach is to bring the supply chain back to a new stable state after the change in the market consumption. Concretely, this mechanism divides the supply chain into five successive states, each of which occurs at different times along the supply chain, i.e., the initial state occurs first with the retailer, next with the wholesaler etc., and the last state occurs first with the wholesaler and finally with the retailer. We assume no capacity limit in the following description.

1) *Initial state:* Suppose we start the process and assume that there has been no market consumption variations for a long time. Therefore, the product stream in the supply chain is stable and equal to the market requirement. In fact, if we assume processes are reliable in order to avoid the only cause of the bullwhip effect in Table I that may apply in this scenario, companies place the same order each week as in the previous week, and this order corresponds to the market consumption: placed O's are equal to incoming O's (principle 1) and $\Theta = 0$ (principle 2). The supply chain is thus in a stable state, i.e., each time period is the same as the previous one.
2) *Perturbation and reaction of the supply chain:* After the single change in the market consumption, companies have to order more than what they were ordering before the change; this is done with O's that follow market consumption. Since there are ordering and shipping lead times, companies will receive for a short period the quantity of products they were receiving before the abrupt change in market consumption (as if there were an inertia in product flow), and inventory level will thus decrease. This is the reason why companies ask for more products than O to reconstitute their inventory: this state lasts a single placement of order in which nonzero Θ's are sent.

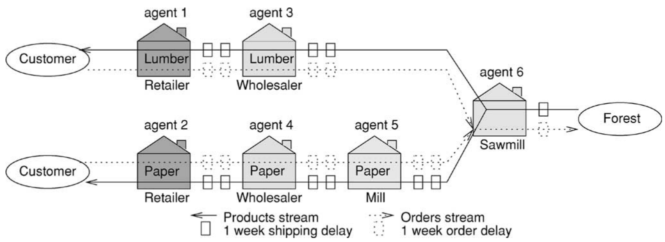

Fig. 4. Model of forest supply chain in the Québec wood supply game.

3) *Wait for the effect of the reaction:* Since O's have increased since the second state, the flow of products received by every company will increase soon. As long as this flow increase has not reached a company, this company's inventory decreases. Notice that no Θ's are emitted.
4) *Stabilization of the supply chain*: This state begins at a company when this company receives products corresponding to the new Os, which stabilizes its inventory. But such a stabilization is too low; it will be so until products corresponding to Θ's and sent by the most upstream company arrive, i.e., until the next state.
5) *Adjustment of the stable state*: The supply chain remains in the fourth state only for some weeks, because Θ's sent in the second state have triggered a larger batch of products to arrive in the company, which puts inventory at a desired level. In fact, when the Θ's sent in the second state have arrived at the most upstream company, a big batch of products was sent down the supply chain to increase all inventories. In other words, this last state is the same as the fourth one, except that inventories are increased to the desired level.

From these five states, we can see that inventory variations last a longer time for retailers than for the most upstream supplier(s), because retailers enter state 2 first and the state 5 last. In fact, inventory stabilization comes from upstream suppliers and travels down to retailers in this model. As a consequence, if managers only consider the duration of backorders, *retailers are more disturbed than suppliers*. Since this is one of the main contributions of this paper, we will come back to this point later.

IV. Experimental Validation

The experimental validation of our two principles is made on the model of the QWSG, which is first described. We also introduce seven ordering schemes IS0-1, IS0-2, IS0-3, IS1, IS1 + P, IS2, and IS2 + P, where only IS1 + P, and IS2 + P follow our two principles. Then, the experimental results are described for one particular market consumption pattern, and finally for 18 other patterns.

A. QWSG

The QWSG is a classroom exercise that simulates the material and information flows in a production-distribution system, as illustrated by Fig. 4, and was designed to make players aware of the bullwhip effect. Compared to the classic beer game that has been used to study supply chain dynamics, the QWSG has a divergent product flow to increase its relevance to the Québec forest sector. Note that the main difference between the original "wood supply game" (the "father" of the QWSG from Fjeld [6], [7], and the "son" of the beer game from Sterman [8]) and our QWSG is in the length of the lumber and paper chain which is either the same (Fjeld's games) or different (our game); this change is due to differences between North European and Québec wood industries.

TABLE II
Experimented Ordering Schemes

Scheme	O	Θ
IS0-1	S minus inventory level when inventory is lower than s	zero
IS0-2	incoming O minus inventory variation, when positive	zero
IS0-3	incoming O plus λ times order variation, when positive	zero
IS1	incoming O	incoming Θ minus inventory variation
IS1+P	incoming O	incoming Θ plus λ times order variation
IS2	customer consumption minus inventory variation, when positive	zero
IS2+P	customer consumption	incoming Θ plus λ times customer consumption variation

Fig. 4 shows how six players (human or software agents) play the QWSG. The game is played by turns: each turn represents a week in reality and is played over five days; these five days are played in parallel by each player. On the first day, the players receive their inventory (these products were sent two weeks earlier by their supplier, because there is a two-week shipping lead time), and advance shipping between suppliers and their customers. On the second day, the 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. On the third day, the players record their inventory or backorders. On the fourth day, the players advance the order slips. On the last day, the players place an order with their supplier(s) and record this order. To decide what to order, the players compare their incoming orders with their inventory/backorder level. This decision is made in our simulation by applying one of the seven ordering schemes ISx and ISx + P in Table II, where ISx means information sharing at level x and +P that our two principles are satisfied. To describe these schemes, we call $[Op_w, \Theta p_w]$ the order $[O, \Theta]$ placed in week w, $[Oi_w, \Theta i_w]$ the incoming $[O, \Theta]$ in the same week, $[Ob_w, \Theta b_w]$ any $[O, \Theta]$ backordered before w, D_w the market consumption (which correspond to Oi_w for a retailer), I_w the inventory level, and λ a parameter ruling the emission of Θ. With these notations, the schemes in Table II are as follows.

- *Scheme* IS0-1 *(no information sharing 1):* IS0-1 is a classic (s, S) ordering policy, i.e., when inventory I is lower than s, the company orders for S-I products so that the inventory increases up to S. With the above notations, we have $[Op_w, \Theta p_w] = [S - I_w, 0]$ when $I_w < s$ and $= [0, 0]$

TABLE III
MARKET CONSUMPTIONS PATTERNS

Week	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	...	49	50
1. Step	11	11	11	11	17	17	17	17	17	17	17	17	17	17	17	17	...	17	17
2. Inv. step	17	17	17	17	11	11	11	11	11	11	11	11	11	11	11	11	...	11	11
3. Dirac	11	11	11	11	17	11	11	11	11	11	11	11	11	11	11	11	...	11	11
4. Inv. Dirac	17	17	17	17	11	17	17	17	17	17	17	17	17	17	17	17	...	17	17
5. Increase	11	11	11	11	12	13	14	15	16	17	18	19	20	21	22	23	...	56	57
6. Decrease	57	57	57	57	56	55	54	53	52	51	50	49	48	47	46	45	...	12	11
7. Weak seas.	11	11	11	11	12	11	10	11	12	11	10	11	12	11	10	11	...	12	11
8. Medium seas.	11	11	11	11	12	13	12	11	10	9	10	11	12	13	12	11	...	10	9
9. Strong seas.	11	11	11	11	12	13	14	13	12	11	10	9	8	9	10	11	...	8	9
10A→J Uniform	Ten uniform integer random distribution on [eleven, seventeen]																		

else, with $(S, s) = (Oi_w, 0)$, (cf. [**?**] for explanations of this choice).

- *Scheme* IS0-2 *(no information sharing 2):* Companies that use Scheme IS0-2 do not know the market consumption because neither $[O, \Theta]$ nor information centralization are used. Each order is a unique number calculated by subtracting the inventory variation to the client's order; when this value is negative, nothing is ordered. More formally, the orders placed with IS0-2 are $[Op_w, \Theta p_w] = [Oi_w + (I_{w-1} - I_w) + (Ob_w - Ob_{w-1}), 0]$ when $Oi_w + (I_{w-1} - I_w) + (Ob_w - Ob_{w-1}) \geq 0$, and $= [0, 0]$ otherwise.
- *Scheme* IS0-3 *(no information sharing 3):* Scheme IS0-3 resembles IS0-2, except that each order is calculated by subtracting λ times the order variation of the client's order. In our experiments, this calculation always gives a positive result. The main difference between IS0-2 and IS0-3 is that IS0-2 is based on order and product flows, while IS0-3 is only based on order flow. Now, $[Op_w, \Theta p_w] = [Oi_w - \lambda(Oi_{w-1} - Oi_w), 0]$ if $Oi_w - \lambda(Oi_{w-1} - Oi_w) \geq 0$, and $= [0, 0]$ otherwise.
- *Scheme* IS1 *(information shared with direct suppliers):* Similarly to IS1 + P, IS1 uses $[O, \Theta]$ orders, but Θ's now depend on inventory level. There is the same relation between IS1 and IS1 + P, as between IS0-2 and IS0-3. Here, $[Op_w, \Theta p_w] = [Oi_w, \Theta i_w + (I_{w-1} - I_w) + (Ob_w - Ob_{w-1})]$ when $\Theta i_w + (I_{w-1} - I_w) + (Ob_w - Ob_{w-1}) \geq 0$, and $= [Oi_w, 0]$ otherwise.
- *Scheme* IS1 + P *(information shared with direct suppliers + our two principles):* IS1 + P is the first ordering scheme based on our two principles. Companies use the lot-for-lot ordering pattern in O's and use Θ's to manage their inventory: client's Θ's are transmitted to the supplier, and λ times the order variation represents the inventory variation that must be balanced by Θs. In these conditions, $[Op_w, \Theta p_w] = [Oi_w, \Theta i_w - \lambda(Oi_{w-1} - Oi_w]$.
- *Scheme* IS2 *(information centralization):* Scheme IS2 is similar to IS0-2, except that information centralization is now used. Companies base their orders on the actual market consumption instead of on client's orders, in order to react quicker to changes of end-customer demand: $[Op_w, \Theta p_w] = [D_w + (I_{w-1} - I_w) + (Ob_w - Ob_{w-1}), 0]$.
- *Scheme* IS2 + P *(information centralization + our two principles):* IS2 + P is an improvement on IS1 + P, because information centralization speeds up the multicast of market consumption information, while both schemes satisfy our two principles. Companies base their orders on the actual market consumption instead of on client's orders: $[Op_w, \Theta p_w] = [D_w, \Theta i_w - \lambda(D_{w-1} - D_w)]$.

The application of one of these seven schemes to place an order is done on the last day of each week. Thereafter, the game continues with a new day 1, and so on. Each position is played in the same way, except the Sawmill: this position receives two orders (one from the LumberWholesaler, the other from the PaperMill) that have to be aggregated when placing an order to the forest. The Sawmill can evaluate its order by basing it on the lumber demand or on the paper demand. In the following experiments, the Sawmill places an order equal to the mean of these two possible orders (in particular, we will see that this may cause backorders in Figs. 6(d) and 7(d), because inventory levels do not stabilize on their initial level, which forces us to use nonempty initial inventory for the Sawmill in these two figures). Moreover, the Sawmill receives one type of product, and each unit of this product generates two units: a lumber unit and a paper unit. That is, each incoming unit is split in two: one piece goes to the Sawmill's lumber inventory, and the other goes to its paper inventory.

When we add information sharing $[O, \Theta]$ to this game, players have to manage the new piece of information Θ. When players receive Θ's with incoming orders O's in the second day of a week, they process them as actual orders. Since the incoming O's change (they always change when $\Theta \neq 0$), players add a positive quantity to Θ's before transmitting them upstream. In IS1 + P and IS2 + P, this addition depends on the parameter λ, whose setting will be presented in Section IV-B.

We now present the behavior of the QWSG under these seven schemes. We first detail the simulation outputs of the Step demand, in order to verify whether the flows in the supply chain stabilize when the market consumption is steady after a single change. The Step demand is presented in Table III in which we can see that both (lumber and paper) markets buy 11 products during the first four weeks, followed by 17 products until the end of the simulation. Table III presents the demand over 50 weeks, which is the duration (i.e., one year) of our simulations. Note that Uniform demands 10A to 10J are ten instances of an integer random distribution on [eleven], [seventeen] that are generated once for all simulations. In contrast to the first nine patterns, the demand in the 10 Uniform patterns may be different in the two markets (e.g., 12 lumber units and 17 paper units

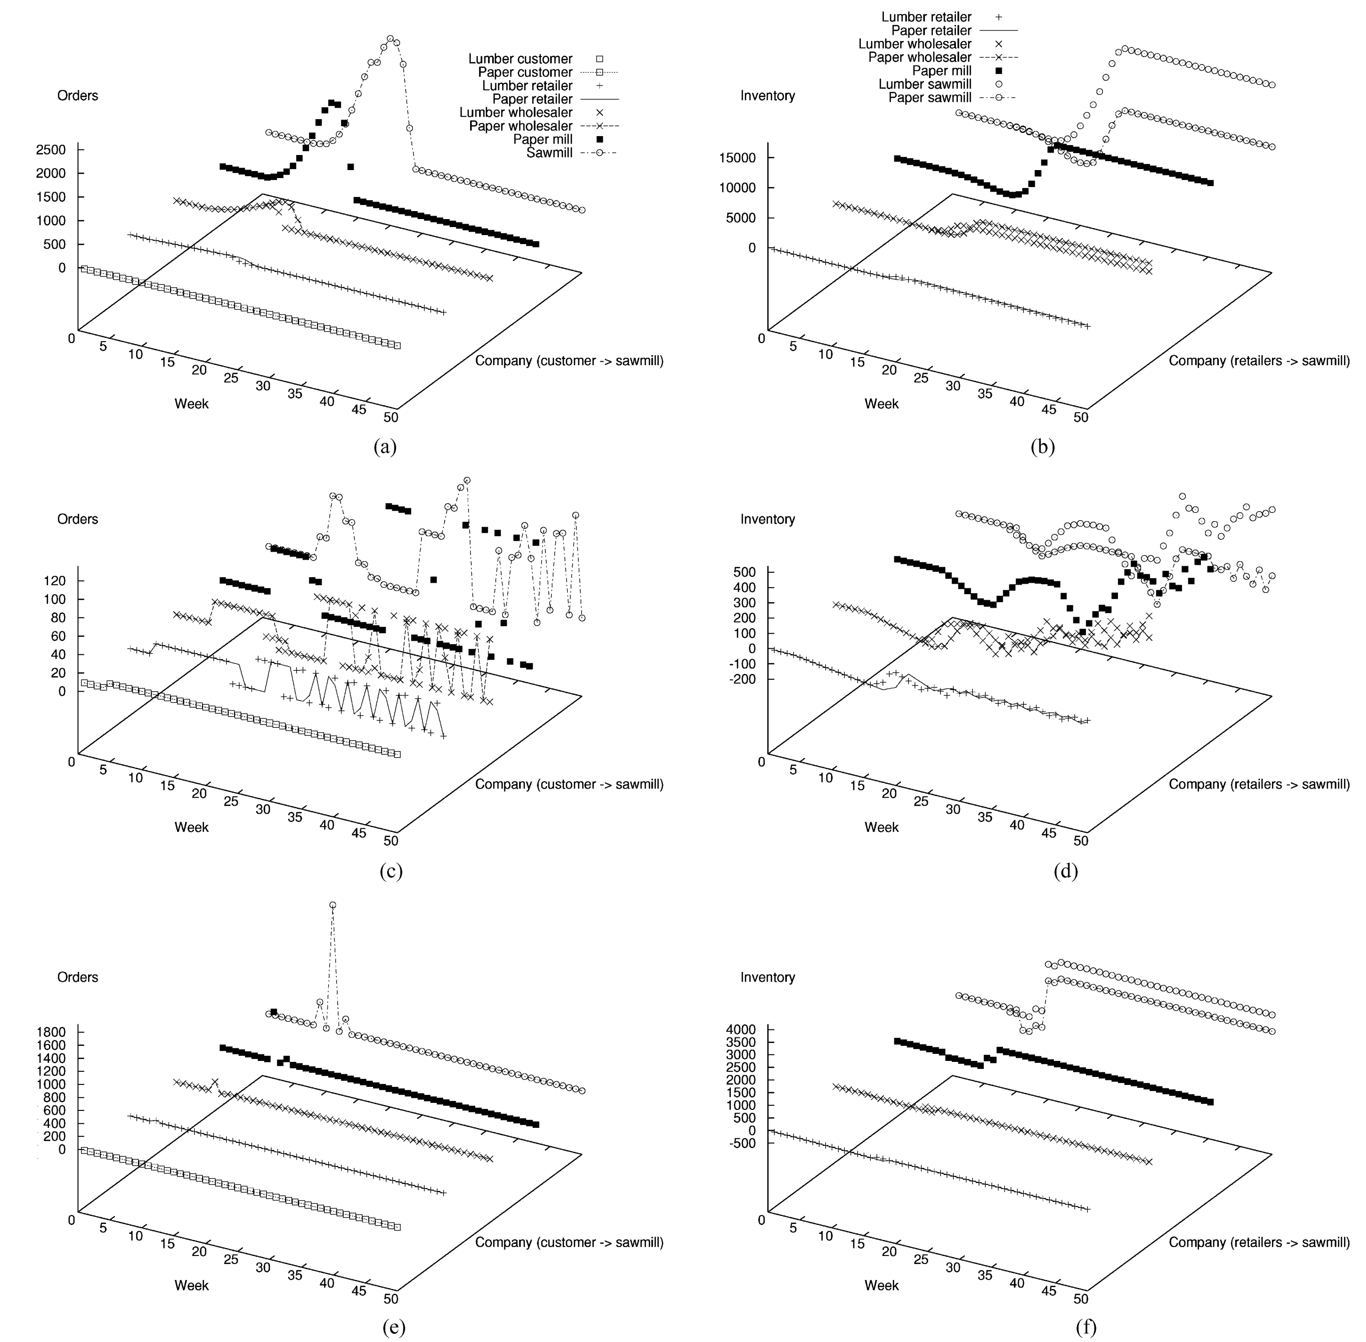

Fig. 5. Results with the three ordering schemes without information sharing. (a) Orders with IS0-1. (b) Inventories with IS-1. (c) Orders with IS0-2. (d) Inventories with IS0-2. (e) Orders with IS0-3. (f) Inventories with IS0-3.

are bought during the same week), which we have not detailed in this paper. Then, we will aggregate the simulation outputs of the 18 other demand patterns of Table III with four metrics in Section IV-C.

B. Experimental Results With the Step Demand Pattern

We now detail the behavior of the QWSG when the market consumption is a Step pattern. Figs. 5(a), 5(c), 5(e), 6(a), 6(c), 7(a), and 7(c) present the demands in the supply chain. In these figures, the first two curves from the bottom (they are one over the other) represent the consumption of both markets under the Step demand. Next, the two second curves show LumberRetailer's and PaperRetailer's placed orders. The two third curves indicate that of LumberWholesaler and PaperWholesaler. Finally, the next-to-last curve represents PaperMill's orders, and the last curve represents Sawmill's orders. Similarly, Figs. 5(b), 5(d), 5(f), 6(b), 6(d), 7(b), and 7(d) present inventory levels (and backorders when inventory level is negative) in the supply chain. In these figures, the first two curves from the bottom show LumberRetailer's and PaperRetailer's inventory level. The second two curves

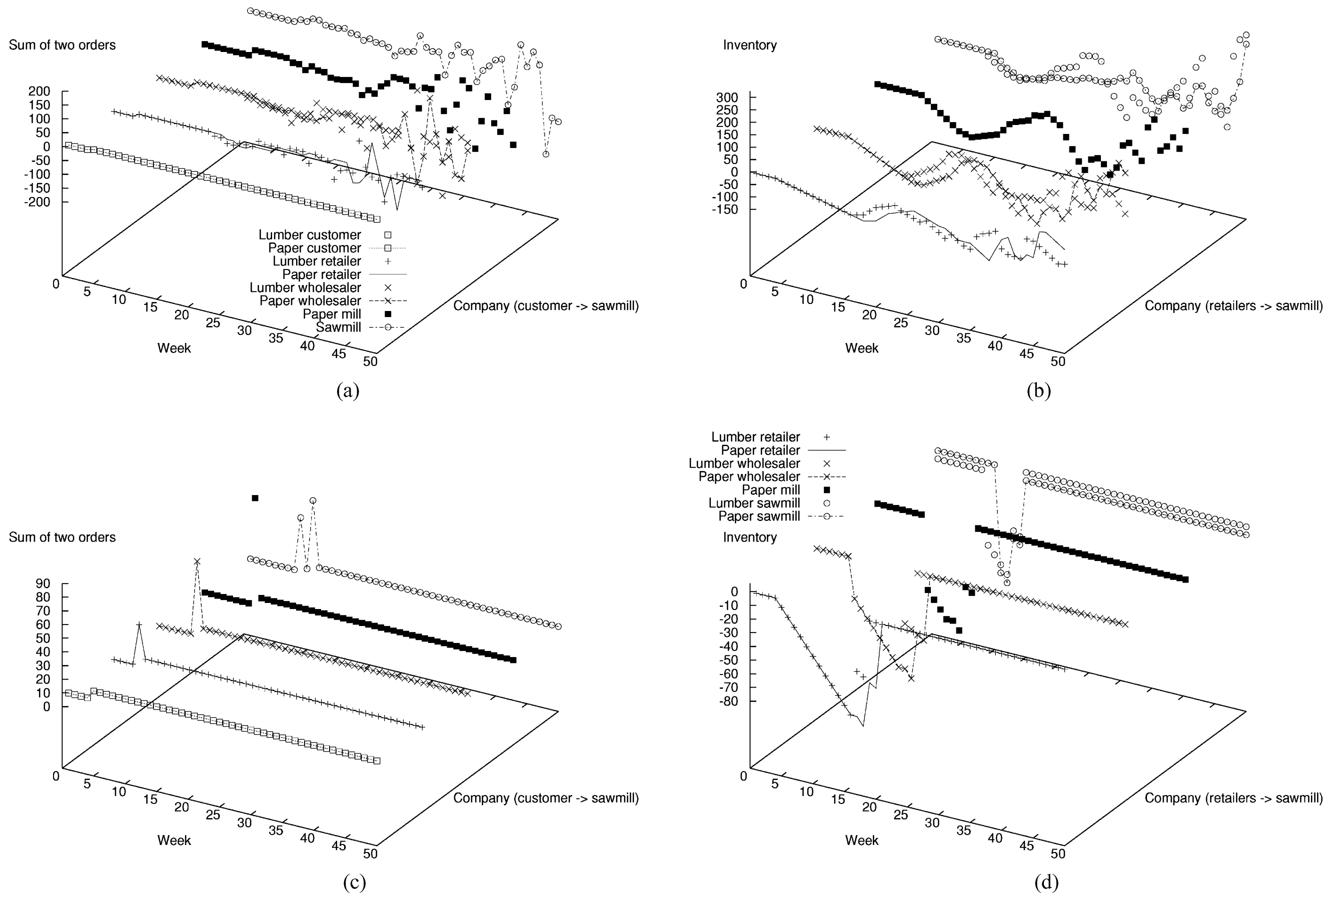

Fig. 6. Results with the two ordering schemes in which information is shared with direct suppliers. (a) Orders with IS1. (b) Inventories with IS1. (c) Orders with IS1 + P. (d) Inventories with IS1 + P.

indicate LumberWholesaler's and PaperWholesaler's inventory. Finally, the next-to-last curve represents PaperMill's, and the last two curves are Sawmill's lumber and paper inventories.

Note that the ranges in Figs. 5–7 are adapted to the simulation outcomes. For this reason, small fluctuations may appear large, because of the scale. Orders and inventories are measured in simulated units, as in QWSG, and not in real units. Moreover, when $[O, \Theta]$ orders are used, Figs. 6(a), 6(c), and 7(c) represent the sum of the two types of orders $(O + \Theta)$, and Figs. 6(b), 6(d), and 7(d) represent the sum of inventory and backordered O's and Θ's. Finally, note that these 14 (7 for orders + 7 for inventories) figures are obtained with empty initial inventories, except for the Sawmill with schemes IS1 + P and IS2 + P ($\forall i \in \{1, 2, 3, 4, 5, 6\text{-lumber}\}, I_1^i = 0$ with all schemes, $I_1^{6-\text{paper}} = 0$ with IS0-1, IS0-2, IS0-3, IS1 and IS2, and $I_1^{6-\text{paper}} = 6$ with IS1 + P and IS2 + P), as placing a single order covering the lumber and paper requirements of the Sawmill leads to less "pretty" curves due to backorders. However, in the next section, Tables IV–VII will be obtained with empty initial inventories for all companies ($\forall i, I_1^i = 0$ in Section IV-C).

In general, orders and inventories fluctuate more, i.e., the bullwhip effect is greater and inventories more poorly managed, with Schemes IS0-1, IS0-2, and IS0-3 in Fig. 5, than with Schemes IS1 and IS1 + P in Fig. 6, while the latter are worse than Schemes IS2 and IS2 + P in Fig. 7. If we look more in detail, we can notice the two following insights. First, orders never stabilize with IS1, while they do with IS1 + P. This shows an important point: order stabilization, i.e., the reduction of the bullwhip effect, is not only due to information sharing with $[O, \Theta]$ orders, but also to the two proposed principles, which should necessarily be satisfied *together*. The second insight deals with the factor λ ruling Θ. We can check that IS1 + P and IS2 + P make the supply chain behave as described in Section III-C, except that the "reaction" (emission of Θ's) is achieved in two different ways: either one emission of Θ's with IS2 + P as soon as the market consumption changes, or several late emissions with IS1 + P. These emissions of Θ's correspond to the peaks of orders in Figs. 6(c) and 7(c). These two different behaviors make two types of position appear. The first type of position in the supply chain concerns the retailers with either ordering scheme IS1 + P or IS2 + P, as well as the other companies with IS1 + P. In fact, when the market consumption changes, such companies: 1) ship more product to fulfill their demand; 2) order in O the new value of the market consumption; and 3) order in Θ additional products to adjust their inventory to its initial level. The second type of position deals with any company using IS2 + P (except retailers, whose case belongs to the first case), because it only has to carry out 2) and 3) immediately, since 1) will only come later. As a result, the first type of position has $\lambda = 4$ and the second type $\lambda = 2$. Since Θ's encapsulate the bullwhip effect and are ruled by λ, a *greater* bullwhip effect occurs in the first type of position.

TABLE IV
STANDARD DEVIATION OF PAPER MILL'S ORDERS

	Market consumption	Scheme						
		IS0-1	IS0-2	IS0-3	IS1	IS1+P	IS2	IS2+P
1 Step	1.6	547.7	51	105.7	42.7	10.6	7.1	4.5
2 Inversed Step	1.6	345.6	33.2	37.4	28.8	10.6	13.1	4.5
3 Dirac	0.8	90.8	18.7	111.8	17	15.2	4.4	6.2
4 Inversed Dirac	0.8	778.2	25.6	142.8	43.5	15.2	8.1	6.2
5 Increase	14.4	590.8	108.1	22.2	16.6	16.7	16.3	16
6 Decrease	14.4	1,928.2	27.9	20.3	41.7	16.7	20.6	16
7 Weak seasonality	0.7	205.5	3.6	98.4	10.6	11.6	0.9	4.2
8 Medium seasonality	1.2	38.7	25.3	73.7	15.2	11.7	4.7	3.2
9 Strong seasonality	1.8	100.7	27.1	59.3	19.8	11.7	12	6.5
10A Uniform random	1.8	576.1	30.3	275.4	37.7	32.8	10.1	13.6
10B Uniform random	1.7	644.8	39.1	206.4	38.8	25.5	11.2	10.6
10C Uniform random	1.7	43.3	36.1	215.1	41.4	28.1	9.4	11.7
10D Uniform random	1.7	349.2	35	181.8	50.1	26.6	13.5	10.2
10E Uniform random	1.9	64	38.9	244.3	51.9	31.5	9.4	13
10F Uniform random	1.8	125.9	35.3	214.4	37.6	26.8	7.1	11.9
10G Uniform random	1.9	281.8	37.9	246.1	23.1	31.5	6.7	14.7
10H Uniform random	1.9	378.6	39.5	243.6	68.9	30.5	10	11.7
10I Uniform random	1.8	540.1	36.1	263	46.8	31.3	7.3	13
10J Uniform random	1.8	754.5	28.9	261.4	33.5	28.4	8.8	12.9

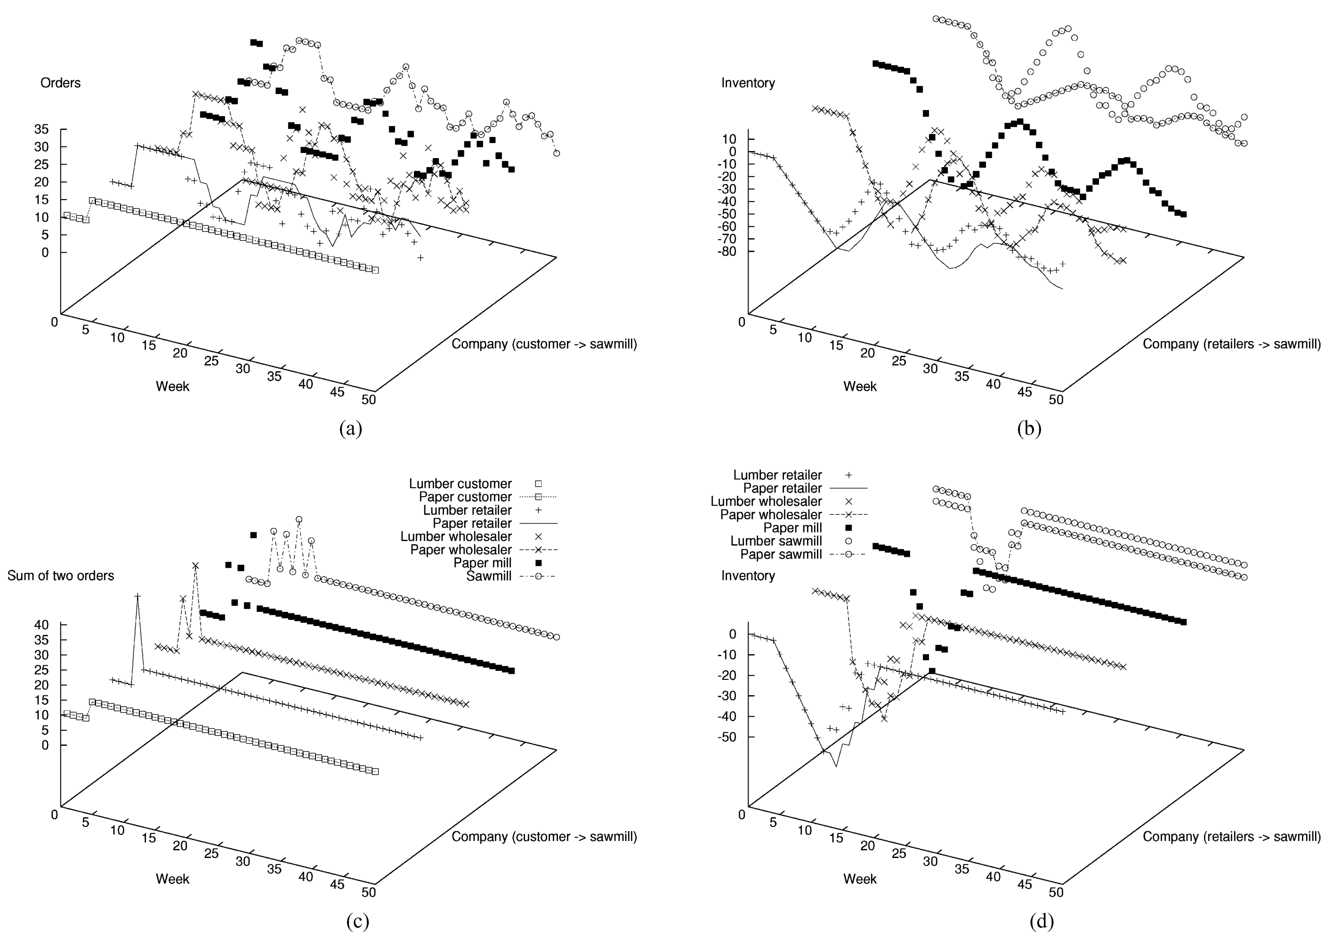

Fig. 7. Results with the two ordering schemes with information centralization. (a) Orders with IS2. (b) Inventories with IS2. (c) Orders with IS2 + P. (d) Inventories with IS2 + P.

C. Experimental Results Under the 19 Demand Patterns

We have executed our seven ordering schemes with the 19 market consumption patterns illustrated in Table III. We do not provide graphs as in Section IV-B in order to save space, but we instead consider four metrics of the efficiency of the seven schemes.

1) Standard Deviation of Orders: This is the direct quantification of the bullwhip effect. Table IV summarizes this data for the PaperMill, e.g., the standard deviation of the orders placed with IS0-1 is 547.7 under the Step demand. This number has to be compared with the market consumption, which has a standard deviation of only 1.6 with this demand.

TABLE V
SUPPLY CHAIN COSTS

	Scheme						
	IS0-1	IS0-2	IS0-3	IS1	IS1+P	IS2	IS2+P
1 Step	1,152,612 $	40,055 $	181,641 $	29,556 $	9,035 $	21,833 $	6,626 $
2 Inversed Step	576,093 $	37,034 $	65,593 $	17,437 $	3,237 $	8,426 $	2,417 $
3 Dirac	192,925 $	9,206 $	226,310 $	6,798 $	1,319 $	2,930 $	989 $
4 Inversed Dirac	1,406,648 $	14,080 $	237,074 $	21,373 $	1,939 $	3,887 $	1,369 $
5 Increase	1,130,332 $	51,681 $	31,071 $	95,308 $	46,754 $	80,770 $	32,620 $
6 Decrease	3,542,566 $	25,022 $	24,111 $	17,888 $	8,754 $	17,436 $	6,207 $
7 Weak seasonality	182,966 $	678 $	244,467 $	2,409 $	2,163 $	279 $	1,194 $
8 Medium season.	76,903 $	12,619 $	140,492 $	7,255 $	7,938 $	3,237 $	3,699 $
9 Strong season.	213,305 $	20,774 $	104,918 $	14,742 $	11,970 $	9,031 $	7,527 $
10A Uniform random	1,107,530 $	19,525 $	666,181 $	12,930 $	9,234 $	5,017 $	5,561 $
10B Uniform random	1,254,093 $	27,959 $	622,179 $	17,929 $	12,140 $	5,711 $	4,684 $
10C Uniform random	158,395 $	26,502 $	751,719 $	26,019 $	12,768 $	8,826 $	6,529 $
10D Uniform random	640,874 $	40,100 $	618,905 $	20,865 $	10,175 $	5,763 $	6,447 $
10E Uniform random	145,021 $	23,916 $	720,748 $	13,937 $	12,144 $	5,017 $	5,574 $
10F Uniform random	280,938 $	20,946 $	727,832 $	12,767 $	8,076 $	5,892 $	4,240 $
10G Uniform random	679,511 $	22,385 $	789,886 $	10,180 $	10,677 $	6,866 $	7,392 $
10H Uniform random	820,017 $	40,087 $	821,050 $	24,549 $	12,009 $	7,973 $	9,286 $
10I Uniform random	1,164,409 $	27,251 $	825,016 $	19,540 $	12,721 $	7,971 $	7,184 $
10J Uniform random	1,568,480 $	20,203 $	701,809 $	13,027 $	10,018 $	6,484 $	5,911 $

TABLE VI
SUM OF PAPER MILL'S BACKORDERS

	Scheme						
	IS0-1	IS0-2	IS0-3	IS1	IS1+P	IS2	IS2+P
1 Step	8,943	1,418	905	747	135	2,081	84
2 Inversed step	5,755	157	65	215	16	157	0
3 Dirac	1,461	246	1,021	186	39	206	45
4 Inversed Dirac	12,942	300	967	473	113	233	91
5 Increase	9,617	3,259	417	664	126	7,956	96
6 Decrease	32,178	48	0	70	0	0	0
7 Weak seasonality	3,129	43	203	154	143	22	48
8 Medium seasonality	608	247	107	333	374	206	81
9 Strong seasonality	1,643	255	189	556	422	651	164
10A Uniform random	9,257	221	233	581	366	297	262
10B Uniform random	10,616	447	247	542	604	345	115
10C Uniform random	531	388	473	511	343	632	158
10D Uniform random	5,433	118	605	593	360	179	104
10E Uniform random	6,127	623	714	505	393	374	175
10F Uniform random	1,996	437	779	259	321	474	132
10G Uniform random	4,999	549	1,323	551	600	630	369
10H Uniform random	6,451	352	772	1,033	367	393	88
10I Uniform random	10,221	807	1,638	966	922	933	287
10J Uniform random	13,596	361	217	229	209	300	104

2) Overall Supply Chain Cost: This is an indirect measure of the bullwhip effect, which is presented in Table V. Of course, this metric is more important for companies than the previous one, because the companies' objective is to maximize their profit rather than reduce the bullwhip effect. Like in QWSG, each unit in inventory costs \$1 per week, and each unit in backorder costs \$2 per week for the Sawmill. For the other companies, we multiply each company's cost by $(1 + 0.37/50)^k$ to represent the increase of the product value due to logistics operations (37% of product value), according to Nahmias [36], where $k = 0$ for the Sawmill, $k = 1$ for the LumberWholesaler and the PaperMill, $k = 2$ for the LumberRetailer and the PaperWholesaler, and $k = 3$ for the PaperRetailer.

3) Sum of Backorders: As presented in Table VI, this is a metric for the customer service level. This measure is included in costs, but it is also interesting to consider it separately. The sum of backorders has to be minimized, because when it is zero, clients have the products they want, or else they have to wait for their availability. Indeed, backorders can be avoided by overstocking, which increases inventory holding costs.

4) Standard Deviation of Inventory Levels: This is used to choose the target inventory level, which has to increase in order to avoid stockouts as the standard deviation of inventory increases. In other words, safety inventory has to be increased when inventory level fluctuates more. Quite the opposite, when an inventory is more stable, its level can be decreased by changing the target inventory level, which reduces companies' costs (i.e., the second metric) without increasing their sum of backorders (i.e., the third metric). Data in Table VII present the fluctuation of PaperMill's inventory; the higher these values are, the more PaperMill's inventory fluctuates, and the best possible value is zero, i.e., always steady inventory and never backorders, but this is attainable only when the whole demand is perfectly known in the future, which is not possible in practice and in our model because there are always forecasting errors.

Note that the greater the numbers in Tables IV–VII are, the worse it is. We now draw general conclusions from these four tables. First, in general, IS0-1, IS0-2, and IS0-3 give worse results than IS1 and IS1 + P, which show that information sharing improves the overall efficiency of a supply chain. Similarly, in general, IS1 and IS1 + P give worse results than IS2 and IS2 + P, which also show that information centralization is better than simple point-to-point information sharing. Second, IS1 + P is almost always more efficient than IS1, which leads us to believe

TABLE VII
STANDARD DEVIATION OF PAPER MILL'S INVENTORIES WHEN BACKORDERS ARE MEASURED AS NEGATIVE INVENTORY LEVELS

	Scheme						
	IS0-1	IS0-2	IS0-3	IS1	IS1+P	IS2	IS2+P
1 Step	4,495.5	180.2	324.7	65.8	22.8	23.9	14.4
2 Inversed step	2,766.7	153.6	108.5	42.2	20.8	24.1	9.4
3 Dirac	714.4	42.6	405.7	19.5	7.8	6.9	6
4 Inversed Dirac	6,370	54.1	441.6	69.2	11.1	12.5	6
5 Increase	4,764.8	185.2	30.8	131.3	42.2	114.3	30.5
6 Decrease	15,538.9	97.5	30.2	42.8	16	50.6	9.6
7 Weak seasonality	1,235.8	2	608.4	6.4	3.1	0.5	1.4
8 Medium seasonality	289	69.9	321.6	13.8	13.8	6.7	8.7
9 Strong seasonality	840.3	98	238.8	44.4	27.2	27.8	17.4
10A Uniform random	4,524.3	72.4	1,714	31.6	21.8	15.2	12.9
10B Uniform random	5,310.2	114.2	1,671	57.1	22.4	16.3	8.3
10C Uniform random	194	110.2	1,979.7	95.5	20.9	17.4	9.7
10D Uniform random	2,713.4	137	1,303.9	63	23.9	23.2	10.7
10E Uniform random	436.1	96	1,894.2	43.9	23.2	11.2	10.4
10F Uniform random	994.8	101.6	1,672.7	51.6	17.4	15.1	10.2
10G Uniform random	2,478.1	107.6	1,771.9	25.6	22.5	8	11.2
10H Uniform random	3,346.4	124.7	2,045.9	69.6	23.1	16.5	11.1
10I Uniform random	4,318.6	134.1	1,787.3	50.2	22.3	10.6	11.3
10J Uniform random	6,753.4	99	1,902.5	42.2	20.9	11	10.4

that our two principles are correct. Unfortunately, it is difficult to generalize this conclusion to the simulations with information centralization. In this case, this conclusion also holds with nonrandomized demands, because IS2 + P is, in general, more efficient than IS2, but results are less obvious with randomized demands (patterns 10A to 10J) in which IS2 is more efficient than IS2 + P in Table IV, as efficient in Table V, and less efficient in Table VI and VII.

V. DISCUSSION

Our experimental results conform to our predictions described in Section III. With IS1 + P and IS2 + P, each company orders exactly its suppliers' order, and over- or underorders *only once* as a reaction to each market consumption change. Since this reaction has to be correctly interpreted by its supplier, information sharing based on $[O, \Theta]$ orders permits this supplier to distinguish between the market change (visible in O) and adjustments triggered by this change (visible in Θ). This stabilizes orders in the whole supply chain at the same level, i.e., at the actual market consumption. If we aggregate O and Θ under the form of $X = O + \Theta$, the level O at which companies have to stabilize their orders is lost for upstream companies, because they do not know which part is required by the market and which part reflects adjustments due to lead times. Sharing information about market consumption is a way to align companies behavior on the same goal: to deliver products to clients. From a more general point of view, companies may use Θ in different ways, for example, either to balance production rejects, or to reduce their inventory by sending negative Θ's which will not be interpreted by suppliers as market consumption changes, etc.

Ordering with $[O, \Theta]$ orders and information centralization are very similar, because both allow each company to know the actual market consumption. In fact, when $[O, \Theta]$ orders are used, O's are equal to market consumption, which is the piece of information multicasted in information centralization. The first difference between these two systems lays in the propagation speed of market consumption information: Θ's without information centralization are as slow as the orders, while with information centralization, each company knows in real time and *instantaneously* the market consumption (retailers multicast the market consumption to the whole supply chain). The second difference between $[O, \Theta]$ orders and information centralization is the fact that information centralization requires a stable supply chain structure to allow retailers to know to whom the market consumption has to be multicast, while $[O, \Theta]$ orders use a more *decentralized* approach.

Instead of sharing information, we could try to smooth orders by placing orders based on forecasts. In our experiments, this would make the stabilization period longer. For example, with the Step market consumption, companies would overorder less than with IS0-2, but this overordering would be longer, and during this period suppliers would not know what the actual market consumption is. However, we think each company should only react once to each change in market consumption (second principle), even if this reaction may be prolonged by forecasts. In other words, smoothing orders with a forecast technique is not incompatible with our two principles: we can use both at the same time.

For industrial practitioners, the lesson of this discussion is that they should apply our two principles. However, additional work is required to define what we mean by a "market consumption change" in our second principle, and how to quantify the emission of Θ's when such a change occurs. We propose in Schemes IS1 + P and IS2 + P to base the quantity of Θ's on the demand variation, but this is not enough in practice to take into account breakage, spoilage and theft, for example.

Finally, one crucial result is the inversion of situation between the current bullwhip effect and the supply chain behavior induced by our two principles. In fact, the bullwhip effect disturbs upstream suppliers (the Forest in QWSG) more than retailers, because orders are more stable near the market and order restabilization comes from the market. On the contrary, with our two principles, there are fewer order fluctuations, and what now disturbs the companies are inventory fluctuations. In short, the most disturbed companies with the bullwhip effect are the ones least disturbed by our two principles. In fact, we stated at the end of Section III-C that upstream suppliers are less disturbed by inventory variations than retailers, because these variations last longer in retailers than in upstream suppliers and inventory restabilization comes from the most upstream company. This is also true when information centralization is used in addition to our two principles.

VI. Conclusion

In this paper, we have dealt with the fluctuations of the ordering streams in supply chains, which is known as the bullwhip effect. More precisely, we have presented why lead times are a cause of the bullwhip effect, and proposed two principles to design coordination mechanisms reducing the impact of this cause. These two principles explain why and how companies should share information in order to reduce such fluctuations when they place orders. Precisely, orders are 2-D vectors $[O, \Theta]$, instead of $X(= O + \Theta)$ where O's and Θ's are hidden. Our first principle rules O by stating that "the lot-for-lot ordering policy eliminates the bullwhip effect, but does not manage inventories," while our second principle rules Θ and states that "companies should react only once to each market consumption change." When these two principles are satisfied, O's represent the market consumption transmitted from clients and Θ's are used by companies to adjust their inventory by ordering more or less than O's. We instantiated these two principles in two schemes, where the first scheme applies only our two principles, while the second scheme uses information centralization in addition, i.e., the multicasting in real time and instantaneously of the market consumption by retailers to the whole supply chain.

We compared these two ordering schemes with five others. These comparisons were carried out by using QWSG, which is an adaptation of the wood supply games (i.e., two games derived from the beer game) to the Québec forest industry in order to teach the bullwhip effect. These experiments demonstrated that it is possible to reduce the bullwhip effect when companies over- or underorder to stabilize their inventory, if companies explain to their supplier(s) with $[O, \Theta]$ orders why they over- or underorder. The main idea is that all companies should know retailers' sales in order to understand if an order change is either due to the market consumption or to the bullwhip effect (i.e., adjustments due to lead times, promotions, lot sizing, etc.). This allows a unique wave of products to travel in the supply chain from the most upstream supplier down to the retailers for each change in the market consumption, instead of a persistence of fluctuations caused by this change.

The two most important contributions of this paper are the presentation of why lead times induce the bullwhip effect, and that, our two principles may disturb retailers more than upstream suppliers, which is in contrast to the bullwhip effect in which retailers are less disturbed by order variation than upstream suppliers. This second contribution depends on how backorders are measured: in duration or in amplitude.

As future work, we plan to address the following questions: How to aggregate market consumption information O when a company has several clients, i.e., what would happen with different supply chain structures? How to relax some assumptions in our model, in particular by adding inventory, production and shipping capacities? What happens when companies do not use the same ordering scheme? In particular, which scheme should be used by each company in order to minimize its individual cost?

Acknowledgment

The authors would like to thank the industrial partners of FOR@C, the Research Consortium in E-business in the forest products industry (Université Laval, Québec City, QC, Canada), and the Natural Sciences and Engineering Research Council of Canada, for supporting this research. They also thank the referees for their valuable suggestions.

References

[1] D. Simchi-Levi, P. Kaminsky, and E. Simchi-Levi, *Designing and Managing the Supply Chain*. New York: McGraw-Hill, 2000.

[2] J. W. Forrester, "Industrial dynamics—A major breakthrough for decision-makers," *Harvard Bus. Rev.*, vol. 36, no. 4, pp. 37–66, Jul./Aug. 1958.

[3] 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," *Manag. Sci.*, vol. 46, no. 3, pp. 436–443, 2000.

[4] H. L. Lee, V. Padmanabhan, and S. Whang, "Information distortion in a supply chain: The bullwhip effect," *Manag. Sci.*, vol. 43, no. 4, pp. 546–558, 1997.

[5] ——, "The bullwhip effect in supply chain," *Sloan Manag. Rev.*, vol. 38, no. 3, pp. 93–102, 1997.

[6] D. E. Fjeld, "The wood supply game as an educational application for simulating dynamics in the forest sector," in *Proc. 2nd World Symp. Logist. For. Sect., Supply Chain Manag. Paper Timber Ind.*, K. Sjostrom and L.-O. Rask, Eds. Växjö, Sweden, 2001, pp. 241–251.

[7] E. Y. Haartveit and D. E. Fjeld, "Experimenting with industrial dynamics in the forest sector—A beer game application," presented at the *Symp. Syst. Models Forestry*, 2002.

[8] J. D. Sterman, "Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment," *Manag. Sci.*, vol. 35, no. 3, pp. 321–339, 1989.

[9] M. Wooldridge, *An Introduction to Multiagent Systems*. Hoboken, NJ: Wiley, 2002.

[10] C. Carlsson and R. Fullér, "Soft computing and the bullwhip effect," *Econ. Complexity*, vol. 2, no. 3, 2000.

[11] D. Taylor, "Measurement and analysis of demand amplification across the supply chain," *Int. J. Logist. Manag.*, vol. 10, no. 2, pp. 55–70, 1999.

[12] C. Daganzo, *A Theory of Supply Chains*. Berlin, Germany: Springer-Verlag, 2003.

[13] J. Dejonckheere, S. M. Disney, M. R. Lambrecht, and D. R. Towill, "Transfer function analysis of forecasting induced bullwhip in supply chains," *Int. J. Product. Econ.*, vol. 78, pp. 133–144, 2002.

[14] J. W. Forrester, *Industrial Dynamics*. Cambridge, MA: MIT Press, 1961.

[15] J. A. Kahn, "Inventories and the volatility of production," *Amer. Econ. Rev.*, vol. 77, no. 4, pp. 667–679, 1987.

[16] H. F. Naish, "Production smoothing in the linear quadratic inventory model," *Econ. J.*, vol. 104, no. 425, pp. 864–875, 1994.

[17] J. Butman, "A pain in the (supply) chain," *Harvard Bus. Rev.*, vol. 80, no. 5, p. 31, 2002.

[18] H. J. Scholl, "Looking across the fence: Comparing findings from SD modeling efforts with those of other modeling techniques," presented at the 19th Int. Conf. Syst. Dyn. Soc., Atlanta, GA, 2001.

[19] B. J. Angerhofer and M. C. Angelides, "System dynamics modelling in supply chain management: Research review," in *Proc. 2000 Winter Simulation Conf.*, 1999, pp. 342–351.

[20] J. U. Min and H. Bjornsson, "Agent-based supply chain management automation," in *Proc. 8th Int. Conf. Comput. Civil Building Eng. (ICCCBE-VIII)*, 2000, pp. 1001–1006.

[21] P. Kelle and A. Milne, "The effect of (s, S) ordering policy on the supply chain," *Int. J. Product. Econ.*, vol. 59, pp. 113–122, 1999.

[22] R. Wilding, "The supply chain complexity triangle—Uncertainty generation in the supply chain," *Int. J. Phys. Distrib. Logist. Manag.*, vol. 28, no. 8, pp. 599–616, 1998.

[23] J. C. Fransoo and M. J. Wouters, "Measuring the bullwhip effect in the supply chain," *Supply Chain Manag.*, vol. 5, no. 2, 2000.

[24] P. Hines and M. Holweg, "Waves, beaches, breakwaters and rip currents—A three-dimensional view of supply chain dynamics," *Int. J. Phys. Distrib. Logist. Manag.*, vol. 31, no. 10, pp. 827–846, 2000.

[25] A. M. Geoffrion and R. Krishnan, "E-business and management science-mutual impacts," *Manag. Sci.*, vol. 49, no. 11, pp. 1445–1456, 2003.

[26] National Science Foundation (2004, March 16). Program announcement 98-168 for advanced computational research., [Online]. Available: http://www.nsf.gov/pubs/1999/nsf98168/nsf98168.htm.

[27] S. K. Yung and C. C. Yang, "A new approach to solve supply chain management problem by integrating multi-agent technology and constraint network," in *Proc. 32nd Hawaii Int. Conf. Syst. Sci.*, 1999, pp. 67–79.

[28] L. Yan, "Impact of lead-time distribution on the bullwhip effect and supply chain performance," in *Proc. Amer. Conf. Inf. Syst.*, 2001.

[29] R. Croson and K. Donohue, "Experimental economics and supply-chain management," *Interfaces*, vol. 32, no. 5, pp. 74–82, 2002.

[30] S. O. Kimbrough, D. Wu, and F. Zhong, "Computers play the beer game: Can artificial agents manage supply chains?," *Decision Support Syst.*, vol. 33, no. 3, pp. 323–333, 2002.

[31] F. Chen and R. Samroengraja, "The stationary beer game," *Product. Oper. Manag.*, vol. 9, no. 1, pp. 19–30, Spriag 2000.

[32] T. Moyaux, B. Chaib-draa, and S. D'Amours, "Agent-based simulation of the amplification of demand variability in a supply chain," in *Proc. 4th Workshop Agent-Based Simul. (ABS4)*, pp. 166–172, Montpellier, France, pp. 166–172, 2003.

[33] ——, "An agent simulation model for the Québec forest supply chain," in *Proc. 8th Int. Workshop on Cooperative Inf. Agents (CIA2004)*, Erfurt, Germany, 2004, pp. 226–241.

[34] ——, "Multi-agent coordination based on tokens: Reduction of the bullwhip effect in a forest supply chain," in *Proc. 2nd Int. Joint Conf. Auton. Agents MultiAgent Syst. (AAMAS)*, Melbourne, Australia, 2003, pp. 670–677.

[35] T. Moyaux, "Design, simulation and analysis of collaborative strategies in multi-agent systems: The case of supply chain management," Ph.D. dissertation, Univ. Laval, QC, Canada, 2004.

[36] S. Nahmias, *Production and Operations Analysis*, 3rd ed. New York: McGraw-Hill, 1997.

Thierry Moyaux received the Engineering degree from the École Polytechnique Universitaire de Marseille, Marseille, France, and the M.Sc. degree in computer-integrated manufacturing and computer science from the Université d'Aix-Marseille 3, Marseille, both in 2000, and the Ph.D. degree in computer science from Université Laval, Quebec City, QC, Canada, in 2004.

He is currently a Research Assistant in the Department of Computer Science at the University of Liverpool, Liverpool, U.K. His research interests include multiagent systems and their application to supply chain management. He has authored several conference and journal papers in these areas.

Brahim Chaib-draa received the Diploma in computer engineering from the Ecole Supérieure d'Electricité (SUPELEC) de Paris, Paris, France, in 1978 and the Ph.D. degree in computer science from the Université du Hainaut-Cambrésis, Valenciennes, France, in 1990.

In 1990, he joined the Computer Science and Software Engineering Department of Laval University, Quebec, QC, Canada, where he is a Professor and Group Leader of the Decision and Actions for Multi-Agent Systems Group. His research interests include agent and multiagent technologies, natural language for interaction, formal systems for agents and multiagent systems, distributed practical reasoning, and real-time and distributed systems. He is the author of several technical publications in these areas. He is on the Editorial Boards of *Concurrent Engineering: Research and Applications (CERA), The International Journal of Grids and Multiagent Systems*, and *Applied intelligence*.

Dr. Chaib-draa is a member of the ACM, the IEEE Computer Society, and the AAAI.

Sophie D'Amours received the Ph.D. degree in industrial engineering from the École Polytechnique of Montréal, Montreal, QC, Canada.

She is the General Director of the Research Consortium FORAC studying supply chain management in the forest product industry and a Full Professor at the Université Laval, Quebec City, QC, Canada. Her research interests relate to supply chain design and planning, distributed planning systems, collaborative planning, and multiagent technology.

Dr. D'Amours holds the Canada Research Chair on supply chain planning.