AGENT-BASED SIMULATION OF THE AMPLIFICATION OF DEMAND VARIABILITY IN A SUPPLY CHAIN

Thierry Moyaux, Brahim Chaib-draa
Département d'Informatique et de Génie Logiciel
Université Laval
Sainte-Foy G1K 7P4 (Québec, Canada)
E-mail: {moyaux, chaib}@iad.ift.ulaval.ca

Sophie D'Amours
Département. de Génie Mécanique
Université Laval
Sainte-Foy G1K 7P4 (Québec, Canada)
E-Mail: Sophie.Damours@gmc.ulaval.ca

Abstract

A supply chain is the set of companies producing or carrying products to customers. In such a supply chain, the bullwhip effect is the amplification of demand variability, that is a distortion in information when this information travels from one end of a supply chain to the other. Inefficiencies which are due to this effect are excessive inventory, poor customer service, ineffective transportations, missed production schedules...

A game called the Beer Game is a widely used classroom exercise for demonstrating the dynamics in a supply chain. We focus on an adaptation of this game to the forest industry: the Québec Wood Supply Game. We have simulated this game in a spreadsheet program: this first implementation is the base of the multi-agent simulation presented in this paper where intelligent agents represent companies. These agents will simulate how companies order, produce and store products.

In this paper, we describe the supply chain model in the Québec Wood Supply Game and how we will make it more realistic with agents. However, we present neither our spreadsheet simulation of this game nor our solution to the bullwhip effect.

1 Introduction

Our work adresses the bullwhip effect in forest supply chains. This effect propagates as follows (cf. Figure 1): the retailer places quite constant orders to its wholesaler, whereas this wholesaler places more variable and unpredictable orders to the manufacturer. Such demand fluctuations cost money due to (i) higher inventory levels (companies stockpile because of the high degree of demand uncertainties and variabilities [Lee et al., 1997b]), (ii) agility reduction

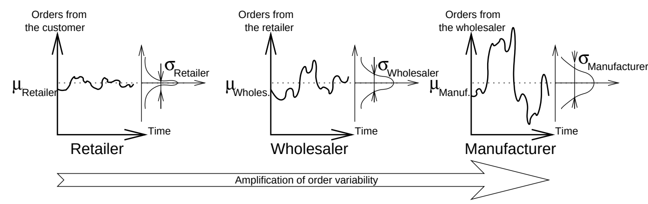

Figure 1: The bullwhip effect [Lee et al., 1997a, Lee et al., 1997b].

(agility is the ability of the supply chain to thrive in a constantly changing, unpredictable business environment [Iaccoca Institute, 1991] in [Rigby et al., 2000]) leading customer service levels to decrease (in our case, customer service is measured by inventory shortages), (iii) ineffective transportations, (iv) missed production schedules...

Our ultimate goal is to improve supply chain efficiency by reducing amplification of demand variability while keeping low inventories and good customer service. To achieve this goal, we have proposed a decentralized coordination mechanism [Moyaux et al., 2003]. We have already studied its performance with a spreadsheet model based on an adaptation of the original Wood Supply Game [Fjeld, 2001, Haartveit and Fjeld, 2002]. This game is an adaptation of the Beer Game [Sterman, 1989] for the forest industry and was designed to make players aware of supply chain dynamics, in particular the bullwhip effect. Our Québec Wood Supply Game (QWSG) take the Québec (a Canadian province) forest industry specificities into account. We describe QWSG game in Section 2.

In this paper, we present a multi-agent simulation designed to study our decentralized coordination mechanism in a more realistic model than QWSG. In fact, original and Québec Wood Supply Games takes into account only one cause of the bullwhip effect (misperceptions of feedback [Sterman, 1989]), while Lee and his colleagues [Lee et al., 1997a], [Lee et al., 1997b] propose four other causes. Our agent-based model is more realistic, therefore more complex, than the QWSG; consequently, it can be run neither by people nor by a spreadsheet program. Therefore, this model is run by intelligent agents. Every agent has the structure of a company based on the SCOR model [Supply Chain Council, 2001]. This multi-agent simulation is presented in Section 3.

2 The Québec Wood Supply Game as a basic model

We now present the QWSG, which is an adaptation to the Québec forest industry of the original Wood Supply Game. Two games, called the "Wood Supply Games" (QWSG is a third variant), were developped [Fjeld, 2001, Haartveit and Fjeld, 2002] based on the structure and dynamics of the Beer Game. Beer and Wood Supply Games are an exercise that simulates the material and information flows in a production-distribution system and were designed to make players aware of the bullwhip effect [Jacobs, 2000]. Compared to the Beer Game [Sterman, 1989] that has been used to study supply chain dynamics, the Wood Supply Game introduces divergent product flows to increase its relevance to the North European forest sector. Our team has adapted this game for the Québec forest sector; we use this version, which is displayed in Figure 2. The main difference between the original Wood Supply Game and our version is in the length of the lumber and paper subchain which is either the same (Fjeld's game) or different (our game). Initial conditions are the same: customer demands, shipping delays and order delays are all equal to four products while each company's inventory contains twelve units of product, and both customers order four products per week for the first four weeks, then eight products until the end of the game. As customer demand is the same as all placed orders, the game begins in a stable state. The change in customer demand in week five is the disturbance which is amplified in the chain (bullwhip effect).

Figure 2 shows how six players (human or software agents) play the game. The game is played by turns: each turn represents a week in reality and is played in parallel by all players for each of the five steps. 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 orders in 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 order to place, players compare their incoming orders with their inventory/backorder level; for example, in our spreadsheet implementation, each player places orders either equal to incoming orders minus inventory variation or equal to customer demand (if we assume every player knows what the market demand is) minus inventory variation. The *decision that affects the bullwhip effect is 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: this position 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 can evaluate its order by basing it on the lumber demand or on the paper demand: in our spreadsheet implementation, the sawmill places an order equal to the maximum of these two possible orders. Moreover, the sawmill receives one type of product from the forest 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.

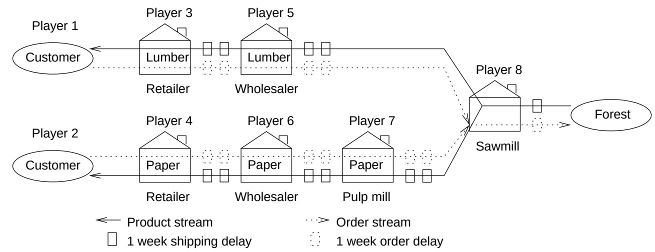

Figure 2: The modelled forest supply chain.

Before we make this model more realistic and we implement it into a multi-agent simulation, we validated it in a spreadsheet program. This first implementation serves in fact as the base for the software of our multi-agent simulation. On the one hand, some elements may be added to the spreadsheet, leading to a more complex implementation. For example, production, storage and transportation capacities can be taken into account. On the other hand, other things can not be added: as there is only one way to discretize time, transportation and ordering delays have to be measured in complete weeks in a spreadsheet instead of in days or in hours. Moreover, when a company places an order, we assumed in the spreadsheet that its supplier always accepts it: we could make this more realistic with the addition of negociation to fix the quantity to ship, the price, the due date... For these reasons, we have to find another way of implementation instead of using a spreadsheet program: we use the multi-agent paradigm.

3 Toward a more detailed model based on the agent paradigm

Many works simulate a supply chain with agents but very few people have focused specifically on the bullwhip effect using multi-agent techniques (e.g. [Kimbrough et al., 2001, Yan, 2001]). The design of our multi-agent simulation of a supply chain is based on our spreadsheet implementation of the QWSG. This game is a very simple model because it only addresses one cause of the bullwhip effect (misperceptions of feedback [Sterman, 1989]), but we can make it more realistic/complex when software agents replace humans by adding other known causes of the bullwhip effect. Lee and his colleagues [Lee et al., 1997a], [Lee et al., 1997b] have proposed four such causes which are seen as the main ones: demand forecast updating, order batching, price fluctuations and rationing and shortage gaming. In this section, we first describe how the QWSG is adapted to agents. Next, we show how each part of a company is modelled in a more realistic way than what is used in the game. Finally, we focus on the implementation of Lee and his colleagues' causes of the bullwhip effect.

3.1 Adaption of the Québec Wood Supply Game to agents

The QWSG is the base of a model that we use for studying the efficiency of different coordination mechanisms. More precisely, we plan to program intelligent agents so that they can play a more sophisticated version of this game. Therefore, the task of every agent will be to simulate a company behavior and to decide when and how much to order. This ordering decision seems very easy to make at first glance, but it is the key problem in the bullwhip effect.

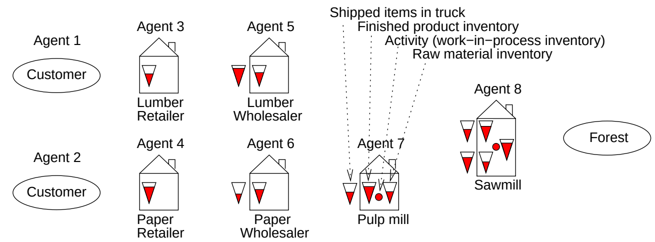

Figure 3: The modelled forest supply chain.

From a global point of view, Figure 3 represents the simulated supply chain. Each company has one or several inventories figured as triangles; the total height of a triangle represents inventory capacity (which is an addition to the QWSG) while its shaded area represents its level. In pulp and saw mills, the circles represent the transformation of a quantity of material (production is another addition to the QWSG): this work-in-process inventory is either full or empty. Both wholesalers are as in the QWSG: they do not have production activity and they have a truck to ship products to their retailer. Other companies differ from the QWSG: both mills manage raw material inventories due to their production activity and neither of the retailers ships products because customers come to buy them. Like in the QWSG, time is discrete[1], each company has product inventory ready to be shipped (Final product inventory I_d) and there are delays between companies (Products in truck I_t). The pulp mill is the only agent that has the basic structure shown in Figures 4 and 5 while the other agents are adapted from these.

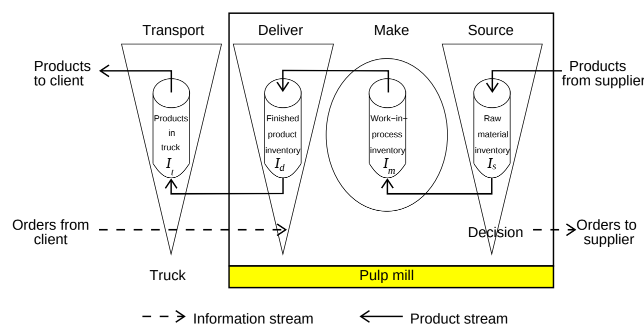

Figure 4: The pulp mill model.

[1] Time discretization makes it easier to ensure that simulation works as we intend it to.

Figure 4 shows the adaptation of the pulp mill in the QWSG to our agent-based simulation. The model of the pulp mill is based on the first level of the SCOR model from the Supply Chain Council in which we add a link Transport between the pulp mill and its client (the paper wholesaler). This link is a truck which takes the place of the two shipping delays in the QWSG. Precisely, these delays are represented in the simulator as an inventory at the visual level and as a queue at the logical level. Product batches go through this queue and are given to the client as soon as they have waited the time which represents the shipping delay. We add capacity to this queue, whereas shipping delays in the Game can ship as many products as needed. We add capacity to inventories too; the inventory in the Game is the Finished product inventory in Figure 4. The QWSG does not make any difference between companies in the distribution network (the retailers and the wholesalers) and production companies (the saw and pulp mills). We make the model more realistic by adding the Make function to the pulp mill, which forces us to add the Source function. Both these functions contain a limited inventory. The work-in-process inventory in the Make function represents the batch being processed, so it is a small inventory which is either full or empty. The Transport, Deliver, Make and Source functions are crossed by the information and the product streams. Figure 5 gives more details about the implementation of these four sets of functions.

3.2 Simulation mechanism

As we have already said, we assume all companies work at the same time and that this time is discrete, i.e. every company waits for the other companies to complete the five steps of the QWSG before begining the next step. Furthermore, there is a bottleneck which is situated at the level of the sawmill, that will allow us to simulate the risks of scarcities (a cause of the bullwhip effect) leading the agents to behave in an opportunist manner. This behaviour is described in the Ordering rule of agents: agents order more than they actually need when there are less incoming products than were ordered.

Roles of the agents

To describe in a general way which agents to implement, we use the Russel and Norvig's PAGE (*Percepts, Actions, Goal, Environment*) description [Russel and Norvig, 1995]. This gives a good idea of what the agents will perceive and do.

- P*ercepts*: client's orders (customer demand is customer's orders).
- A*ctions*: production, inventory and shipment management and order placement.
- G*oals*: ordering in a coordinated way with other agents to reduce the bullwhip effect.
- E*nvironment*: the other agents (the customer is also an agent).

Agents can be divided into several categories:

1. *Customers* are modelled as two agents buying products from their retailers. Customer agents are very simple, because they only place orders according to a distribution law representing different kinds of consumption patterns (constant, constant with little variations, seasonal, increasing, both increasing and seasonnal, etc.).
2. *The distribution network* is composed of the four companies (lumber and paper retailers and wholesalers) that have neither production activity nor raw material inventory to manage: products coming from the supplier do not wait in the raw material inventory I_s because the production capacity is viewed as infinite, so they go directly to the finished product inventory I_d. Moreover, these four agents have to place orders to their suppliers.
3. *The pulp mill* is the basic agent in our model (cf. Figure 5), because it has exactly one of each of the Transport, Deliver, Make and Source components. Compared to the companies belonging to the distribution network, the pulp mill has to manage its production and its raw material inventory I_s. In particular, if the pulp mill produces different kinds of paper, the production has to be planned in order to optimize the production time compared to unproductive times (set-up times...).
4. *The sawmill* works as the pulp mill, except products made in the activity fill two finished product inventories I_d^{lumber} and I_d^{paper} that are then carried by two different trucks I_t^{lumber} and I_t^{paper}.
5. *Forest*: as a most upstream company, the forest has no supplier and therefore does not place any orders. It is assumed to be an infinite source of wood in the QWSG, but we assume in our model that its capacity is given by a cutting plan.

Note: Transportation could have been modelled as truck agents instead of as a function in each company. It would have allowed us to focus on impacts of unanticipated transportation events on the bullwhip effect.

We assume transportation is a problem that can be simply viewed as a queue managed by each company.

Implementation of the agents

Figure 5 represents the software architecture of the pulp mill agent. Eight rules (written in italic in Figure 5) are used to implement this company agent. These rules simulate and improve the five steps per week of the QWSG. These rules are described below according to their role (a general description of the function achieved by the rule), preconditions (what drives the rule) and their action (what the agent does when the precondition is true)[2].

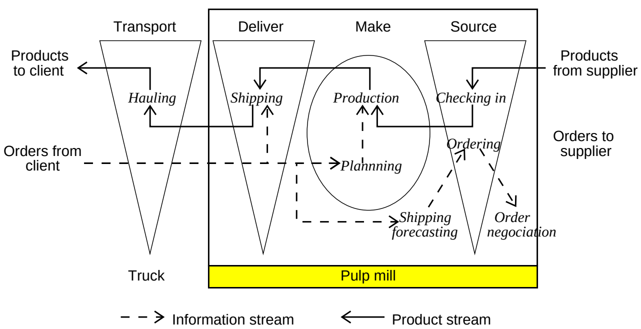

Figure 5: Functions in the pulp mill model.

- *Checking in*

  Role: Receive products coming from the supplier's truck and shipped earlier by the supplier.

  Precondition: arrival of products.

  Action: products carried by the truck are added to the raw material inventory (products are transfered from the upstream I_t to the company's I_s.

- *Production*

  Role: Transform a quantity I_m of raw material into finished products.

  Precondition: Inventory I_m is empty and there is enough products in I_s to fill I_m (I_m is either full or empty).

  Action: A quantity I_m of products is transferred from I_s to I_m and next from I_m to I_d as many times as is possible in the week (several batches can be processed in the week if inventories allow it). If I_d is full, products wait in I_m until there is enough room in I_d.

[2]See Figure 4 for the signification of notations.

- *Shipping*

  Role: Order treatment, that is negociation with the customer orders then shipping these orders to the truck.

  Precondition: an order arrival from the customer or I_d was empty and contains products again.

  Action: The negociation among the agent and its customer is done here with a negociation protocol; in a first approach, shipping does not negociate and ship to I_t products taken in I_d without trying to ship a full truckload. If I_d is not sufficient, shipping stores the missing units as *backorders*.

- *Hauling*

  Role: Send products from truck I_t to the customer's raw material inventory I_s.

  Precondition: the first element of the I_f queue has spent a period equal to the shipping delay.

  Action: products are transferred from I_f to downstream I_s.

- *Planning*

  Role: Plan which product to produce if the company processes different kinds of products.

  Precondition: Each week.

  Implementation: in a first approach, no planning has to be done because only one type of product is processed in the supply chain.

- *Shipping forecasting*

  Role: Anticipate the future demand of the customer.

  Precondition: Each week.

  Implementation: for example, an average of incoming orders in the last three weeks.

- *Ordering*

  Role: Order from the supplier; techniques for reducing the bullwhip effect are implemented here.

  Precondition: Depends on the technique tested for reducing the bullwhip effect. For example, it may be either at the end of each week or each time the inventory I_s is too low.

  Implementation: Depends on the technique tested for reducing the bullwhip effect.

- *Order negociation*

  Role: Negociate the price, quantity, shipping date... with the supplier.

  Precondition: Each time ordering places an order.

Implementation: in a first approach, orders are simply sent to the supplier.

For the moment, we assume companies can not negociate what they order (so the *order negociation* rule is not implemented). Therefore, every agent has three types of messages to manage:

1. *Orders* are messages with fields: *type of product*, *desired quantity* and *desired reception date.*
2. *Shipped batches* are messages composed of many *products* and a *product* is an object with the following fields: *a date of creation* (to calculate the time a product spends in the supply chain), *price* and *type.*
3. *Customer demand* is broadcasted by both retailers to their suppliers. This message is sent at the beginning of each week and only contains what retailers are going to sell during the coming week.

3.3 Adding other causes of the bullwhip effect

All that we have just described will replicate a supply chain. We now focus on a mechanism that takes into account the four causes and solutions of the bullwhip effect proposed by Lee and his colleagues. This mechanism proposes a way to adapt the ordering pattern to Lee et al's causes and solutions. In fact, adding all causes is quite easy, but the adaptation of each company's ordering pattern is more complex. For example, it is easy to create a price fluctuation: the price of an item may be for instance two times the price at which we bought it (to make a profit) minus a tenth of the inventory level (that is, the more products we have, the lower the price we charge). But we must then adapt how companies order when they want to take advantage of sales promotions. So ordering patterns represent how companies behave to take advantage of available opportunities. Moreover, each of the four following causes of the bullwhip effect are simulated one after another, because studying all of them at the same time is not possible. These four causes are [Lee et al., 1997a, Lee et al., 1997b]:

1. *Demand forecasting updating*:

   Proposed solution: make demand data from each company available to its supplier, this allows the supplier to make better forecasts instead of making forecasts on its client's forecasts. This is information sharing.

   Consequence for the agents: a variant of this solution is information centralization at points of sale that broadcast to every company actual market consumption. With this solution, the basic QWSG is adequate enough to study supply chain improvement: each player can base its orders on actual market demand instead of on orders placed by its direct client.

   Example of ordering pattern: incoming orders in the ordering pattern may be replaced by the retailer's incoming orders (i.e. the market consumption).

2. *Order batching*:

   Proposed solution: break up order batches by making transaction costs lower using an electronic commerce system.

   Consequence for agents: production has to be made by batches (that is I_m is either full or empty) and transportation costs depend on the number of trucks used instead of the number of products shipped (so companies will prefer to ship by full truckloads instead of less than full truckloads).

   Example of ordering pattern: placed orders can only take discrete values and these values correspond to full truckloads.

3. *Price fluctuation*:

   Proposed solution: stabilize prices (e.g. EDLP strategy - Every Day Low Price) to avoid customer and company overordering and stockpiling during promotions.

   Consequence for agents: quantity ordered depends on the price of product (client) and this price depends on inventory levels (supplier).

   Example of ordering pattern: placed orders are equal to the needed quantity plus the difference between the nominal price and the current price, where the current price is calculated by the supplier (e.g. the current price is the nominal price plus the difference between the current inventory level and a nominal inventory level).

4. *Rationing and shortage gaming*:

   Proposed solution: eliminate gaming in shortage situations by allocating products to customers in proportion to their past sales instead of in proportion to current orders. Therefore, when a company has several clients but is not able to ship them what they have ordered, the client will not overorder (i.e. gamble) in the hopes of receiving its actual needs.

Consequence for agents: agent overorders if it does not receive what it ordered four weeks ago (because products ordered in week n are received in week $n+4$).

Example of ordering pattern: order the need plus the difference between incoming shipment and what was ordered four weeks earlier.

4 Conclusion

Our goal is to reduce the bullwhip effect, which is the amplification of demand variability. In this paper, we have described how we improve a supply chain simulation (the Québec Wood Supply Game) with intelligent agents in order to make a more realistic simulation of the forest industry. The original Wood Supply Game was created to teach supply chain dynamics and in particular the bullwhip effect. But this game has only one cause of the bullwhip effet: misperception of feedbacks in supply chains. We have simulated the Québec version of this game in a spreadsheet program in order to make a first validation of a solution based on tokens as a coordination mechanism that we propose to the bullwhip effect.

The multi-agent simulation presented in this paper is based on the Québec Wood Supply Chain and its implementation in a spreadsheet program. This multi-agent simulation is designed to validate our token-based mechanism on a more realistic model; in particular, the four causes of the bullwhip effect proposed by Lee and his collegues are added. This leads to a model too complex either to be played by humans or to be implemented in a spreadsheet. The agents running this new model are based on the SCOR model. In particular, we focus on elements relevant to the bullwhip effect, that is those affecting ordering patterns.

5 Acknowledgement

We thank For@c, the research consortium in E-business in the forest products industry (Université Laval, Sainte-Foy, Québec, Canada), for supporting this work.

References

[Fjeld, 2001] Fjeld, D. E. (2001). The wood supply game as an educational application for simulating dynamics in the forest sector. In Sjöström, K. and Rask, L.-O., editors, *Supply Chain Management for Paper and Timber Industries*, pages 241–251, Växjö.

[Haartveit and Fjeld, 2002] Haartveit, E. Y. and Fjeld, D. E. (2002). Experimenting with industrial dynamics in the forest sector - a beer game application. In *Symposium on Systems and Models in Forestry*, Chili.

[Iaccoca Institute, 1991] Iaccoca Institute (1991). 21st century enterprise strategy: An industry led view. Lehig University, Bethlehem, PA.

[Jacobs, 2000] Jacobs, F. R. (2000). Playing the beer distribution game over the internet. *Production and Operations Management*, 9(1):31–39.

[Kimbrough et al., 2001] Kimbrough, S. O., Wu, D., and Zhong, F. (2001). Computers play the beer game: can artificial agents manage supply chains ? *Decision Support Systems.*

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

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

[Moyaux et al., 2003] Moyaux, T., Chaib-draa, B., and D'Amours, S. (2003). 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)*, Hotel Sofitel, Melbourne, Australia.

[Rigby et al., 2000] Rigby, C., Day, M., Forrester, P., and Burnett, J. (2000). Agile supply : rethinking systems thinking, systems practice. 2(3):178–186.

[Russel and Norvig, 1995] Russel, S. and Norvig, P. (1995). *Artificial Intelligence - A modern approch.* Englewood Cliffs, NJ: Prentice-Hall.

[Sterman, 1989] Sterman, J. D. (1989). Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Science*, 35(3):321–339.

[Supply Chain Council, 2001] Supply Chain Council (2001). Scor version 5.0! introduction web cast. http://www.supplychainworld.org/WebCast/SCOR50_overview.ppt.

[Yan, 2001] Yan, L. (2001). Impact of lead-time distribution on the bullwhip effect and supply chain performance. In *Proceedings of Americas Conference on Information Systems.* (Submitted).