The Stationary Beer Game*

Fangruo Chen and Rungson Samroengraja

Graduate School of Business, Columbia University, New York, NY 10027

Phone: 212-854-8694 Fax: 212-316-9180

Booz, Allen & Hamilton Inc., 101 Park Avenue, New York, NY 10178

Phone: 212-551-6568 Fax: 212-551-6815

November 1998; Revised May 1999

Abstract

This paper presents a variant of the popular beer game. We call the new game the *stationary beer game*, which models the material and information flows in a production-distribution channel serving a stationary market where the customer demands in different periods are independent and identically distributed. Different players, who all know the demand distribution, manage the different stages of the channel. Summarizing the initial experience with the stationary beer game, the paper provides compelling reasons why this game is an effective teaching tool.

Keywords: Operations Management, Supply Chain Management, Teaching, Beer Game

---

* To appear in a special issue of *Production and Operations Management* on teaching supply chain management.

1. Introduction

The beer game is an exercise that simulates the material and information flows in a production-distribution system. It has four players: retailer, wholesaler, distributor, and factory. Customer demand (in kegs of beer) arises at the retailer, which replenishes its inventory from the wholesaler, the wholesaler from the distributor, and the distributor from the factory. In each period, the channel members must decide how much, if any, to order from their respective suppliers and the factory must decide how much, if any, to produce. There are transportation leadtimes in shipping the material from one location to another, and there is a production leadtime at the factory. While material flows from upstream to downstream, information flows in the opposite direction through order placements. There is an order processing delay, or information leadtime, between when an order is placed and when the order is received by the supplier. The players share a common objective to optimize the system-wide performance. For more details on the beer game, see, e.g., Sterman (1984, 89).

In the beer game, the customer demand is 4 kegs per period for the first several periods and then changes to 8 kegs per period for the rest of the game. Moreover, the players have no prior knowledge about the demand process. (The numbers 4 and 8 are not important, but the demand pattern and the players' lack of information about it are.) Here in the stationary beer game, the customer demands in different periods are independent and identically distributed, and all the players *a priori* know the demand distribution.

There are compelling reasons why the stationary beer game is an attractive teaching tool. First, it is quite common that companies have some knowledge about the market demand and are able to use that information for planning purposes. This feature is captured in our game since the players know the customer demand distribution. (If one uses the stationary beer game in conjunction with the Barilla SpA (A) case, students can easily see the appropriateness of this assumption.) Second, there exists a theoretical benchmark for the supply chain, i.e., what rational players do and what the optimal supply chain performance is. It is useful to have this piece of information since students often ask what we, the instructors, would have done. (For the original beer game, however, there does not exist such a benchmark.) Third, students often find it an interesting exercise to formulate a replenishment strategy by using the demand distribution, and they are often eager to discuss the rationale behind their strategies. Fourth, the stationary beer game can be used as an example of a periodic-review inventory model that is often taught in Operations Management or Supply Chain Management courses (see, e.g., van Ryzin 1998). The experience with the game helps students visualize important inventory concepts. Finally, with our computer program, it is easy to run the game and to collect and display the results. It is our experience that a single instructor is sufficient for a class of 60 students.

2. The Stationary Beer Game

Consider a supply chain consisting of four stations: a factory, a distribution center, a warehouse, and a retail store. Material flows from upstream to downstream (i.e. from the factory to the distribution center, then to the warehouse, and finally to the retail store), while information in the form of replenishment orders flows in the opposite direction. Both the material and information flows are subject to delays. Different players manage the stations. They only have access to local inventory status and make local replenishment decisions. Customer demand arises at the retail store only. The demands in different periods are independent and identically distributed random variables. The demand distribution is known to all the players. Holding costs are incurred at each station for their on-hand inventories, and backorder costs are incurred only at the retail store for customer backorders. The goal is to minimize the *total* holding and backorder costs incurred in the entire supply chain.

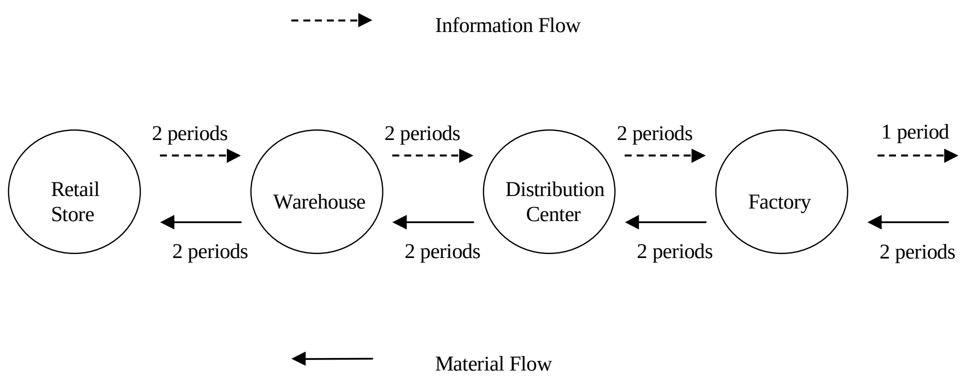

Figure 1. The Supply Chain and Its Leadtime Parameters.

Figure 1 depicts the material and information flows in the supply chain. The figure also specifies the two kinds of delays at each station. Here is an example. Suppose the retail store orders 10 kegs from the warehouse on Monday. The warehouse receives this order on Wednesday. This information delay is due to the administrative steps in processing an order. On Wednesday, however, the warehouse only has 5 kegs of beer, so it ships 5 kegs to the retail store and backlogs the remaining 5. This shipment of 5 kegs arrives at the retail store on Friday. This delay is due to transportation.

The demands in different periods are independent, identically distributed, normal random variables. The normal distribution is discretized and truncated at zero to avoid negative demand values. The demand in a period has mean 50 and standard deviation 20.

Holding costs are assessed at every station for on-hand inventories. Each station must satisfy the orders from its downstream player (or the customers) as much as possible. In case of a stockout, the excess is backlogged. A penalty cost is assessed only at the retail store for customer backorders. This reflects the supply chain's desire to provide good customer service. (In contrast, the original beer game charges a penalty cost at each station. As we will see later in Section 5, when the stations are managed as cost centers, a penalty cost is charged at each station.) Table 1 summarizes the cost parameters.

Location	**Holding Cost ($/keg·period)**	**Penalty Cost ($/keg·period)**
Factory	0.25	0
Distribution Center	0.50	0
Warehouse	0.75	0
Retail Store	1.00	10

Table 1. Cost Parameters.

A computer program is available for playing the game, see the appendix.

3. Theoretical Benchmarks

For the stationary beer game (under the objective of minimizing the long-run average total cost in the supply chain), the optimal strategy for each station is to place orders so as to keep its installation stock at a constant target level, i.e., to follow an installation base-stock policy. The installation stock for a station is its on-hand inventory minus its backorders plus its outstanding orders (i.e., orders placed but not yet received). Notice that the installation stock at a station is local information and thus an installation base-stock policy is feasible under the game's information structure. The optimal target levels for the retail store, the warehouse, the distribution center, and the factory are respectively 280, 225, 214, and 153. The resulting minimum long-run average total cost is $125 per period. (We refer the reader to Chen (1996) for a proof of the optimality of this policy and an algorithm for finding the optimal target levels.)

Figure 2 depicts the replenishment orders and net inventories at each station if all the players follow the optimal strategy. The average cost over the first 30 periods of the game (the default game length) is $196 per period. (This is different from the theoretical benchmark because of the initial conditions of the game. As we increase the length of the simulation, the average cost per period converges to $125.) Notice that the order stream at each station is identical to the customer demand process shifted in time (except for the first few periods due to, again, the initial conditions of the game). Therefore the variance of replenishment orders does not increase from downstream to upstream.

If we eliminate the information delays, the theoretical minimum supply-chain cost drops to $51 per period, a 60% savings. This may result from the use of advanced information technology. Without information delays, the optimal strategy for each player is still an installation base-stock policy, but the target levels are now 157, 117, 110, and 107 for the retail store, the warehouse, the distribution center, and the factory respectively.

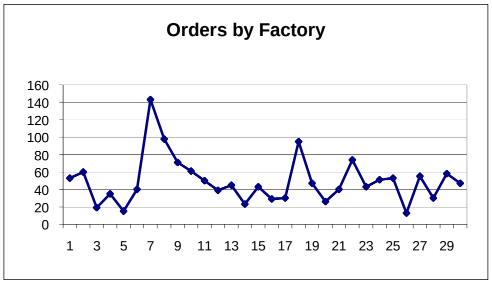

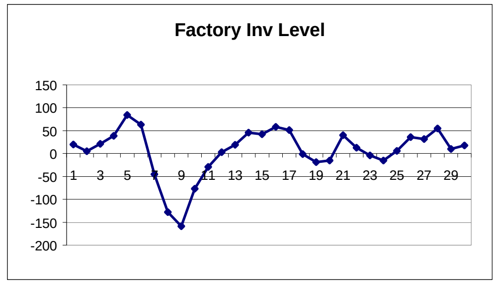

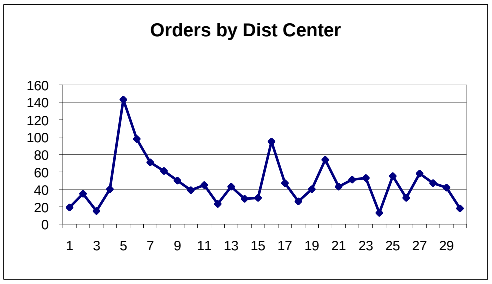

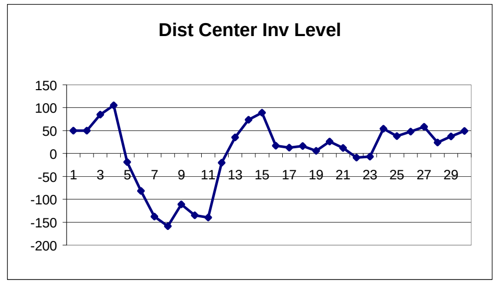

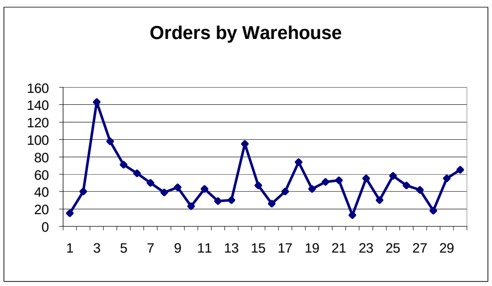

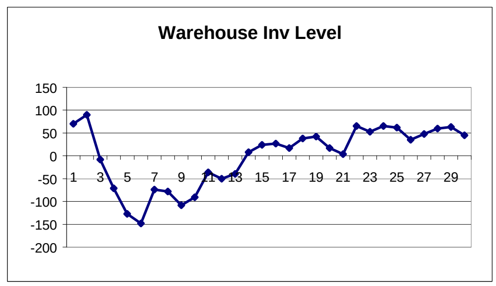

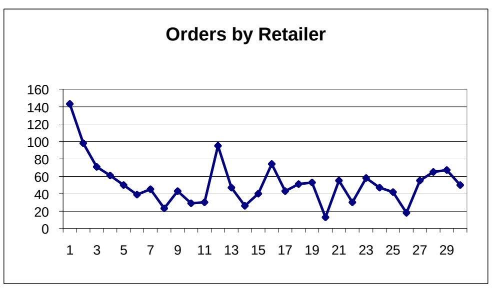

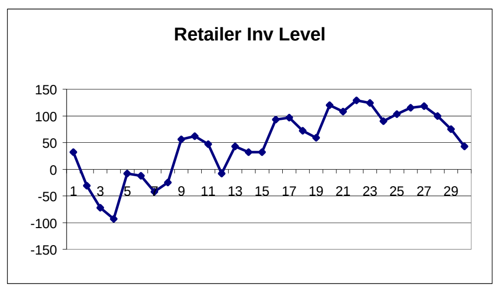

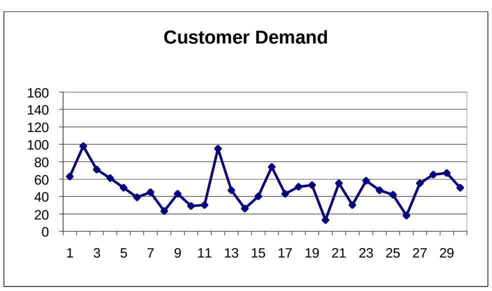

Cost Performance	
Factory	165
Distribution Center	443
Warehouse	625
Retail Store	4,660
Total	5,893
Number of Periods	30
Cost/Period	**196**

Figure 2. Results under the Optimal Strategy.

4. Our Experience

The stationary beer game has now been used in several master-level courses in both engineering and business schools. It has received great reviews from students. Below, we concentrate on our most recent experience in the core Operations Management course here at the Columbia Business School. In the spring semester of 1999, there were eight sections of Operations Management with a total of 500 students. We used the game at the beginning of the Supply Chain Management module of the course. We spent one class (80 minutes) to introduce the game with a demo and to let students play a training version of the game (which the computer program provides). The students then played the game "for real" outside of class and handed in a disk containing the team's results. We spent the following class to debrief the game together with a discussion of the Barilla case. Before the demo session, we asked students to download the computer program from the course's web site, to read the case that describes the game in detail, and to form groups. (We allowed more than one student at each station. Typically, two students managed a station.) The students then brought their laptop computers to the demo session. Each instructor was able to run the demo session without any help from teaching assistants.

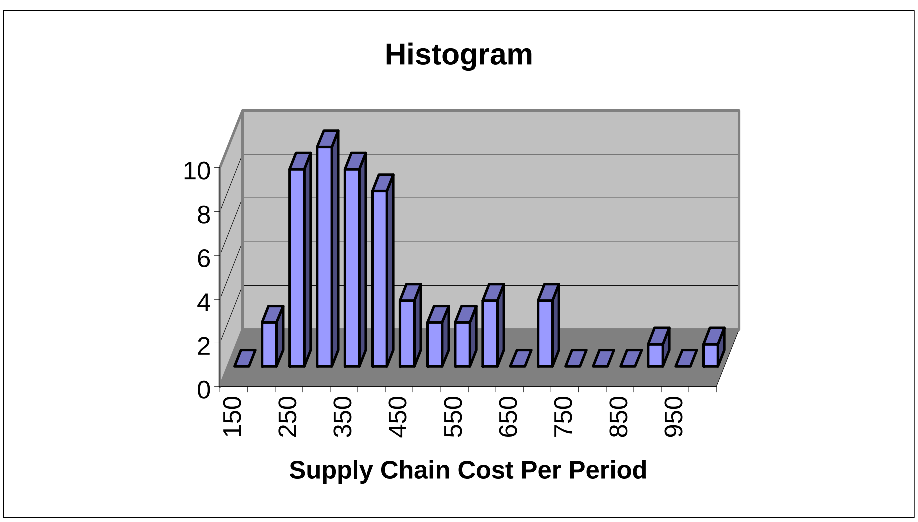

Figure 3. Cost Data from Spring 1999.

Although the supply chain cost is $196 per period (over the first 30 periods) under the (long-run average) optimal strategy, the actual costs achieved by our students varied greatly. Figure 3 is the histogram of the cost data. (After removing those teams with incomplete results, we have 55 teams left. One team achieved a per period cost of $182, which is actually lower than the benchmark cost of $196. This is possible since the optimal strategy minimizes the long-run average cost.)

One of the key observations from the original beer game is the so-called variance amplification phenomenon, i.e., upstream orders tend to be more volatile than the downstream ones. Can the same be said for the stationary beer game? We computed the standard deviation of the order stream for each station in each team. Figure 4 plots the average standard deviation across teams by position. Interestingly, the variance-amplification phenomenon persists. Figure 5 is typical, where orders tend to become more volatile as they go upstream.

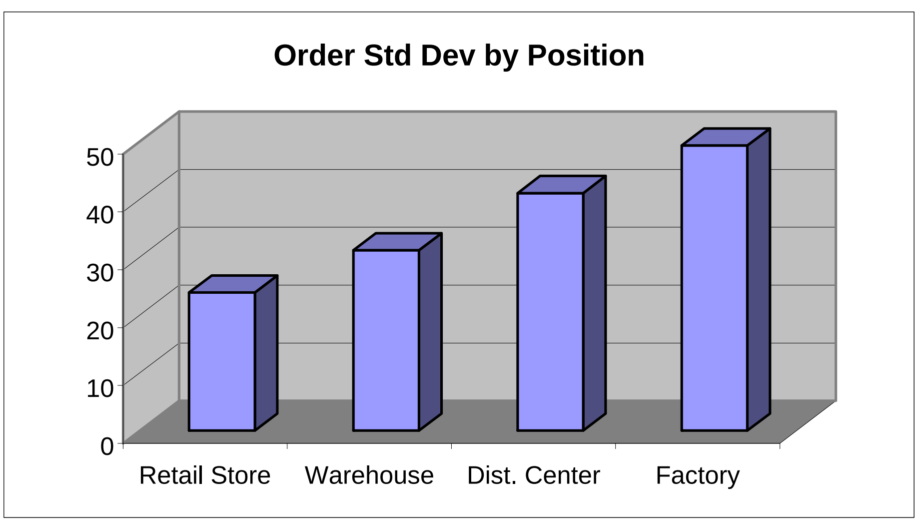

Figure 4. Variance Amplification in Spring 1999.

It is not surprising that the students used the demand distribution to formulate their strategies. For example, consider Figure 6. For each of periods 10 and 11, the warehouse ordered 400 kegs of beer. Apparently, these unusually large orders (relative to the mean demand in a period) did not move the distribution center, which continued to order the mean demand every period. Therefore, the information about the demand distribution helped mitigate, but not eliminate (as Figure 4 shows), the amplification of order variance. (One would wonder what the factory was doing given a steady stream of orders from the distribution center.)

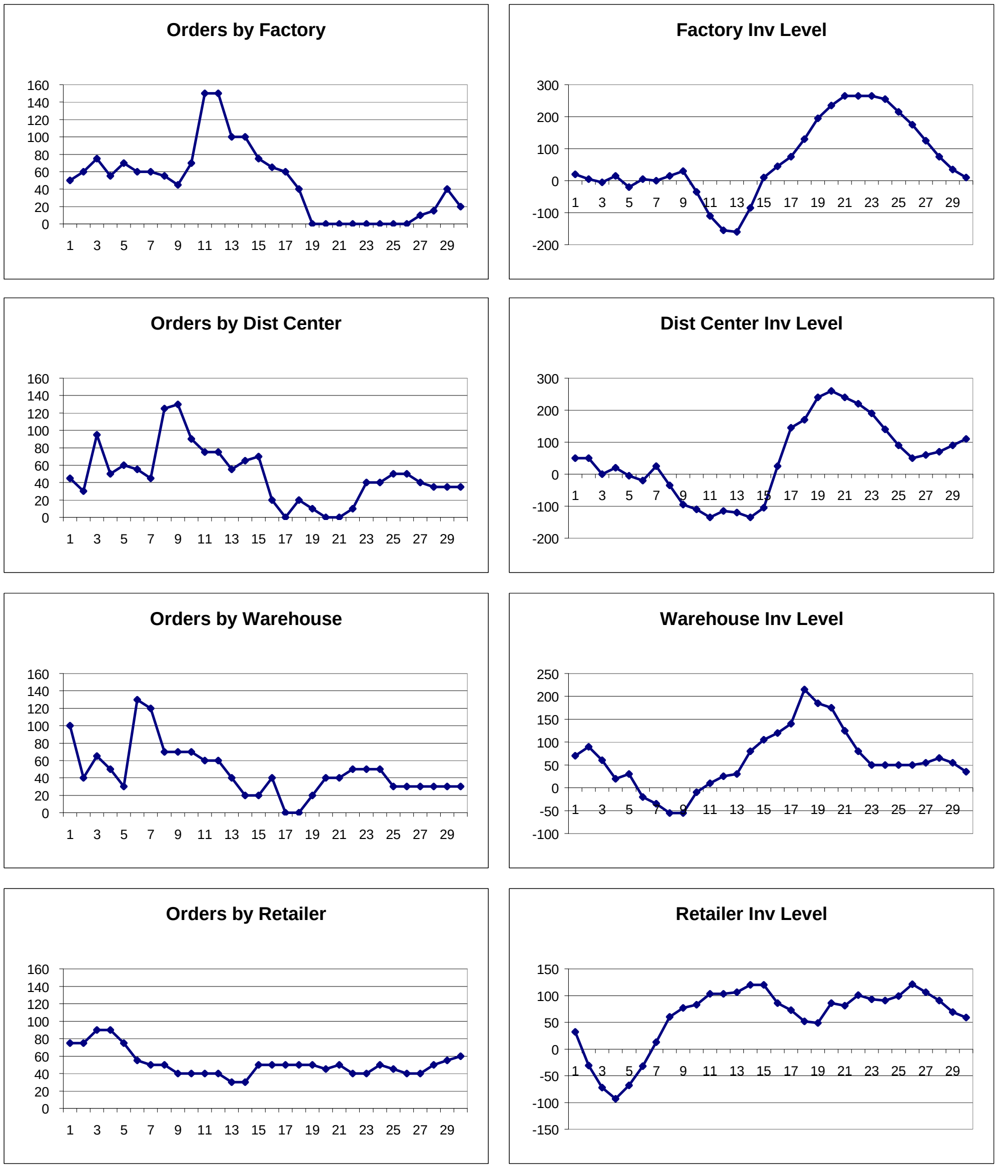

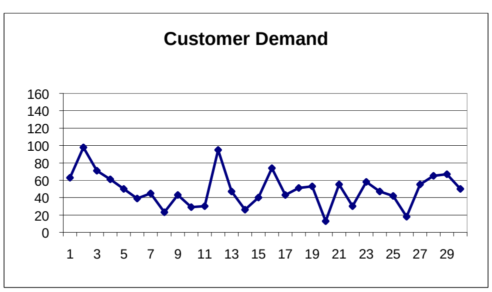

Cost Performance	
Factory	616
Distribution Center	1,123
Warehouse	1,478
Retail Store	5,034
Total	8,250
Number of Periods	30
Cost/Period	**275**

Figure 5. A Typical Example with Variance Amplification.

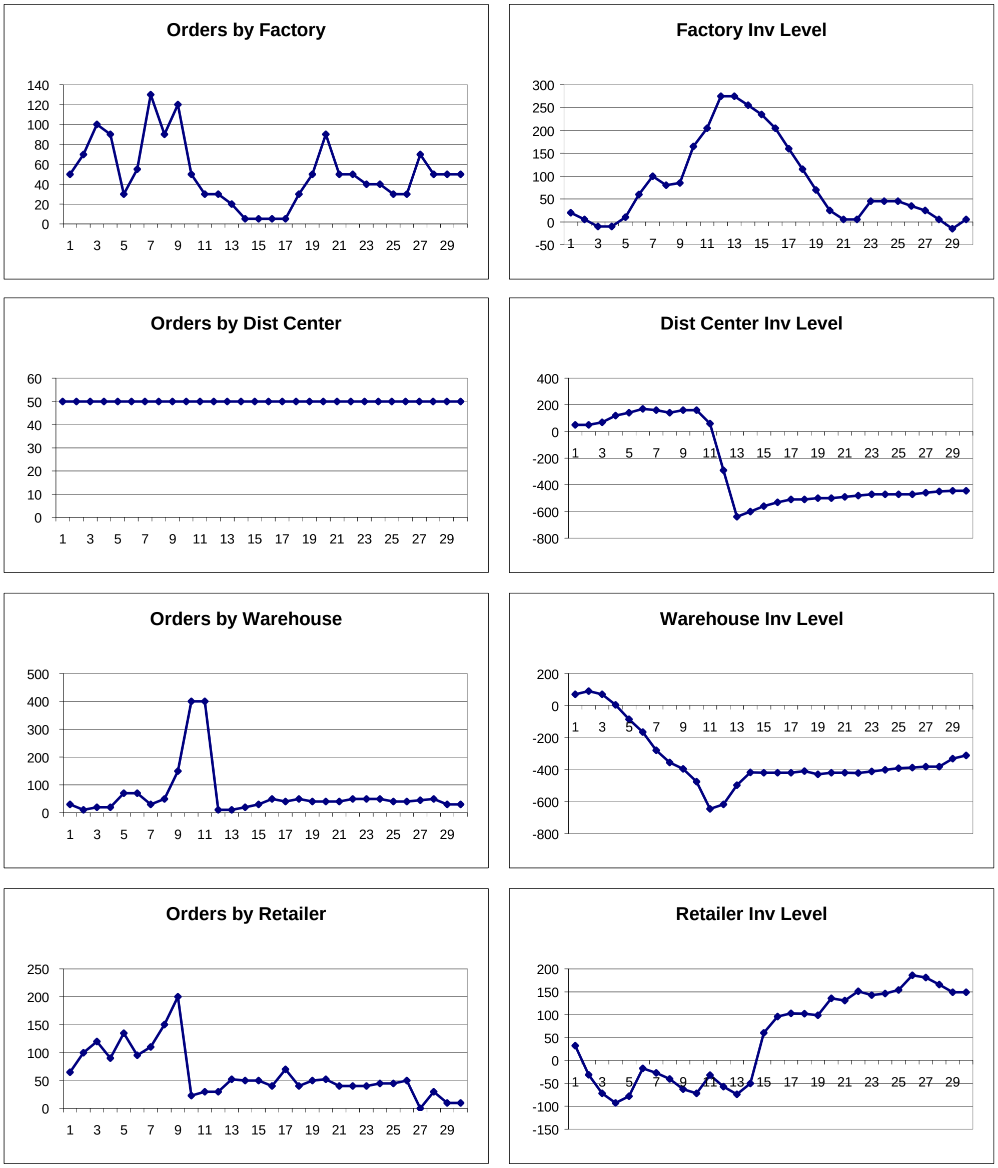

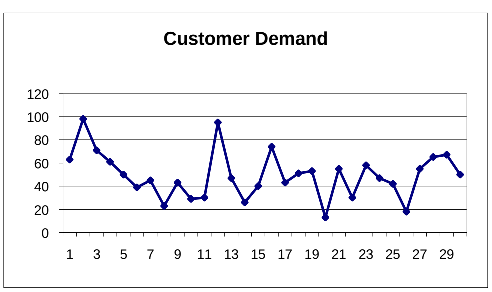

Cost Performance

Factory	640
Distribution Center	640
Warehouse	176
Retail Store	9,244
Total	10,700
Number of Periods	30
Cost/Period	**357**

Figure 6. An Interesting Strategy.

Here, a natural question arises: If each station orders the mean demand (i.e., 50 kegs) every period, what will the supply chain's total cost be? This can lead to a very interesting discussion. After a while, it is not difficult for the students to see that this is a terrible strategy in the long run. The right way is to "manage" the inventory position (net on-hand inventory plus outstanding orders), i.e., placing an order every period so as to maintain a target inventory position. This is a key concept in inventory management. (A station's inventory position is also referred to as its installation stock in inventory theory.) Earlier, we mentioned that the optimal base-stock level for the retail store is $S = 280$. The total leadtime at the retail store is 4 periods (2 periods of transportation leadtime plus 2 periods of information leadtime). Therefore, the total demand during this leadtime has mean $\mu_L = 4 * 50 = 200$ and standard deviation $\sigma_L = \sqrt{4} * 20 = 40$. From the formula $S = \mu_L + z\sigma_L$ we have $z = 2$, implying a fill rate of 99% (see, e.g., van Ryzin 1998).

5. Extensions

The stationary beer game can be easily modified to simulate a supply chain with a different incentive and/or information structure. An alternative incentive structure is to manage the supply chain as cost centers, i.e., each player is charged the costs incurred in his or her own station under a predetermined accounting rule. On the other hand, an alternative information structure is to transmit the customer demand information to the upstream stations. Here is how. When the retail store places an order, it is also required to state what the most recent customer demand was. This demand information is then "tagged" to the order and travels to the warehouse; when the warehouse receives this information, it is required to tag the information to its next order to the distribution center, and so on. Therefore, each upstream player has two pieces of information coming from downstream in each period: a replenishment order and the corresponding customer demand value. The two incentive structures and the two information structures lead to four games (Table 2). Note that Game I here is the stationary beer game.

	Demand Information Not Transmitted	Demand Information Transmitted
Team	Game I	Game III
Cost Centers	Game II	Game IV

Table 2. Four Games.

When the supply chain is managed as cost centers, the costs charged to each station are based on its accounting inventory level, which is determined by assuming that the immediate upstream supplier is perfectly reliable (i.e., it never runs out of stock). Therefore, it is likely that a station's accounting inventory level is different from its actual inventory level. For Games II and IV, each station is charged a holding cost if its accounting inventory level is positive and a penalty cost otherwise. Table 3 provides the accounting costs for each station. These cost parameters have been chosen so that if the players act rationally, the (actual) supply chain cost as measured by using the (actual) cost rates in Table 1 is minimized. (The actual supply chain cost may be different from the sum of the accounting costs.)

Location	**Holding Cost ($/keg·period)**	**Penalty Cost ($/keg·period)**
Factory	0.25	0.3
Distribution Center	0.25	0.4
Warehouse	0.25	0.7
Retail Store	0.25	10.7

Table 3. Accounting Costs for Games II and IV.

The incentive structure determines who the winners are. For Games I and III (with the team structure), the winning team is the one with the lowest supply chain cost. For Games II and IV (with the cost-centers structure), there are winning individuals, instead. Here, the accounting cost at, e.g. the retail store, is compared with the accounting costs of all the other retail stores. The one with the lowest accounting cost is a winner. Therefore, there will be four winners, one for each station. Of course, they may not come from the same team.

When the players are rational, the actual supply chain cost should be identical across the four games. In reality, this may not be the case and the discrepancy can be used to illustrate the impact of accurate demand information and decentralized decision making. The computer program can be used to play all the four games. We hope that Games II, III, and IV will be used in classrooms in the near future.

Acknowledgement: This paper is dedicated to the students in the MBA core course B6801 Operations Management in the spring semester of 1999 at the Columbia Business School. Their input to the development and improvement of the stationary beer game is invaluable. We are grateful to Professors Nelson Fraiman, Fred Silverman, and Garrett van Ryzin for agreeing to try the new game in the OM course and for their helpful suggestions and comments. We would also like to thank the editors of this special issue and the referees of this paper for their suggestions and comments that have significantly improved the exposition of the paper. Financial support from the National Science Foundation, the Columbia Business School, and the Eugene Lang Foundation is gratefully acknowledged.

Appendix: A Computer Program for Playing the Games

Supply Chain Simulations, a Visual Basic for Excel application, was developed for the sole purpose of playing the games described in this paper. (To obtain a copy of the program and related materials, please send an email to the first author at fc26@columbia.edu.) With the program, the games are easy to run and fun to play. The only requirement is that each station have a computer with Excel for Office ’97 (Version 8). Therefore, there are four computers in a team. These computers are stand-alone. (We are in the process of developing a web-based version.) Consequently, the passing of information between two stations is done manually. Below, we present the benefits of using the program from the standpoints of both the players and the instructor together with some screenshots of games in progress.

Benefits to Players:

As anyone who has played the original beer game knows, a significant amount of time is spent executing the mechanics of material and information flows. It is not easy to handle shipments of dozens of pennies (if not hundreds) while the instructor is trying to keep all teams in sync. *Supply Chain Simulations* automates almost all of the game mechanics. Furthermore, since the program acts as an “instructor” guiding the players through every step of the game, confusion is minimized and teams can proceed at their own pace.

The interface was designed to help the students visualize the material and information flows. The screen may be different depending on the game played. Figure 7 and Figure 8 are screenshots from Games I and IV, respectively.

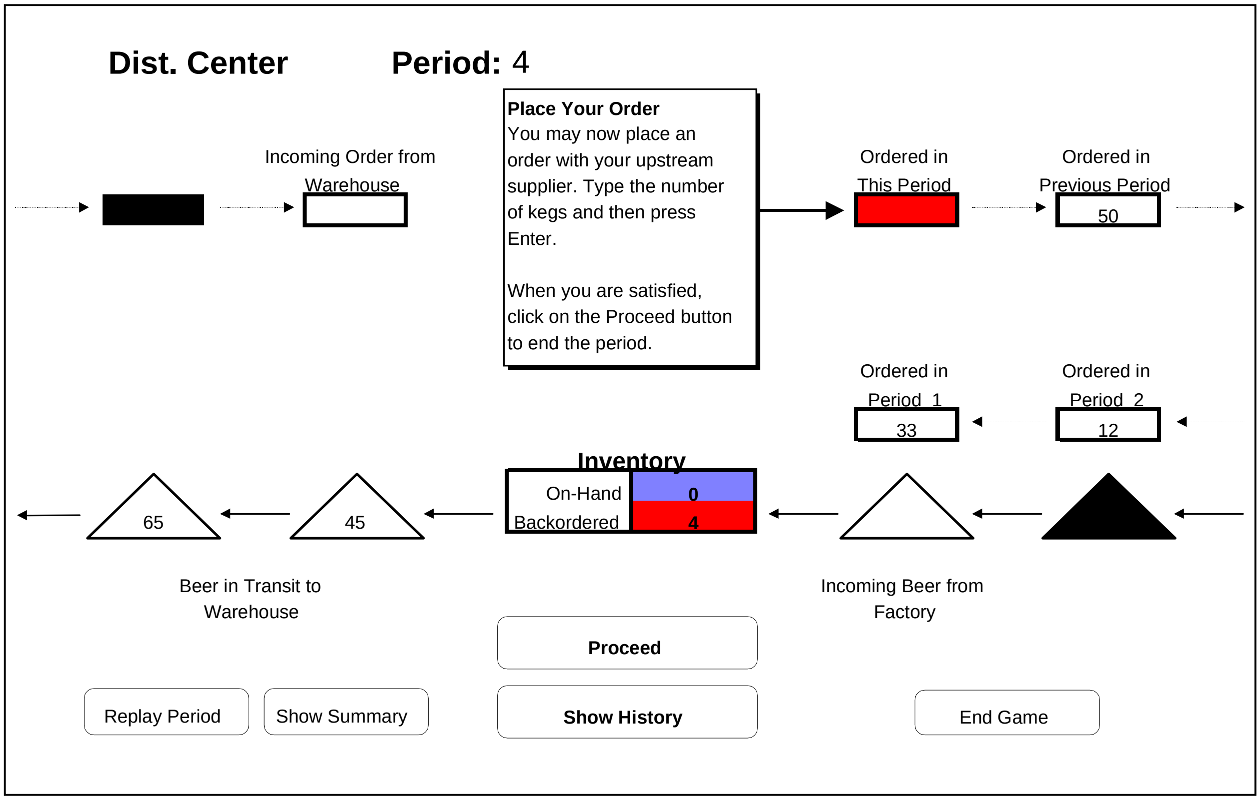

Figure 7. The Main Screen for the Distributor (Game I).

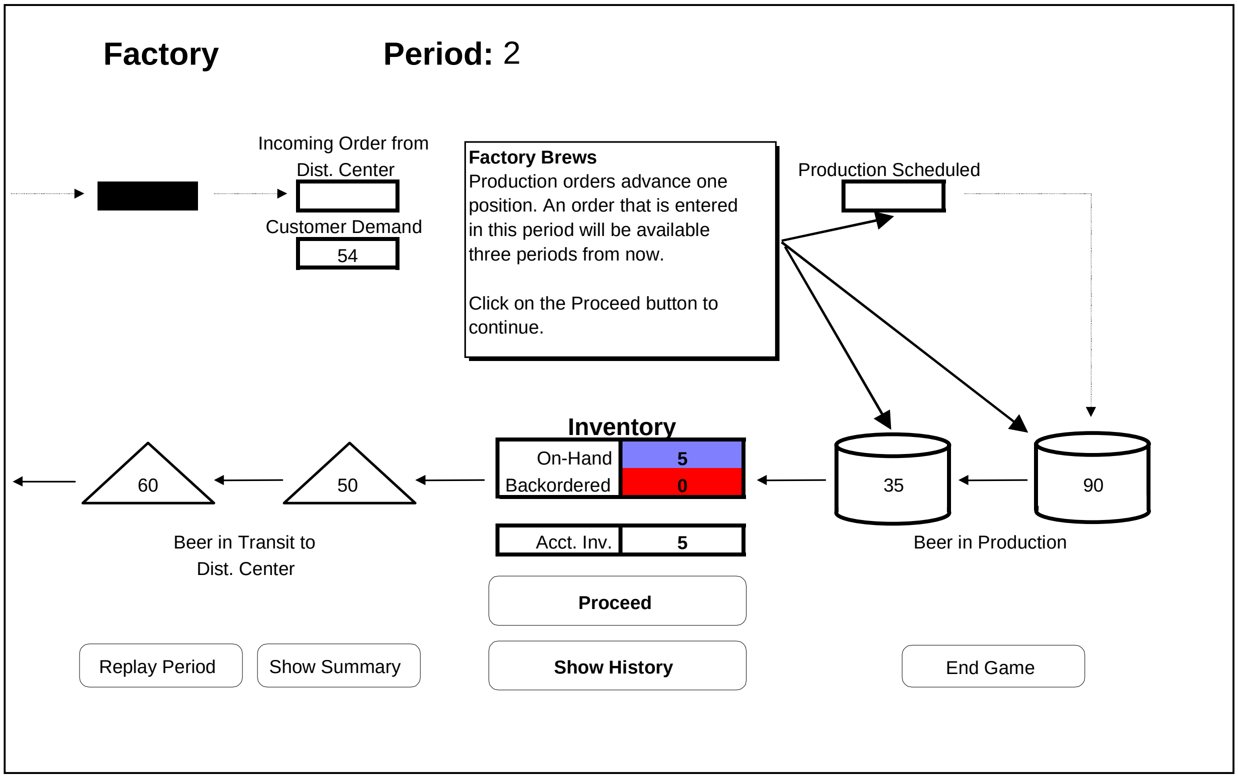

Figure 8. The Main Screen for the Factory (Game IV).

The program tracks all relevant state variables for the players such as inventory level and outstanding orders and presents them in both numerical and graphical formats, see Figure 9 and Figure 10.

Beer History Sheet

Return To Main Screen | Show Graphs | Use the PgUp and PgDn keys to scroll

Position: **Dist. Center**

Period	Starting Net Inven.	Quantity Received	Warehouse Order	Quantity Shipped	Ending Net Inven.	Starting Outstanding Orders	Order to Factory	Ending Outstanding Orders	Holding & Backorder Cost	Cumulative Cost
1	50	24	33	33	41	137	33	146	21	21
2	41	55	44	44	52	146	12	103	26	47
3	52	16	65	65	3	103	50	137	2	48
4	3	42	49	45	(4)	137	60	155	0	48
5	(4)	20	42	20	(26)	155	65	200	0	48
6	(26)	15	32	15	(43)	200	65	250	0	48
7	(43)	50	40	50	(33)	250	65	265	0	48
8	(33)	65	12	45	20	265	40	240	10	58
9	20	65	30	30	55	240	30	205	28	86

Figure 9. The History Screen.

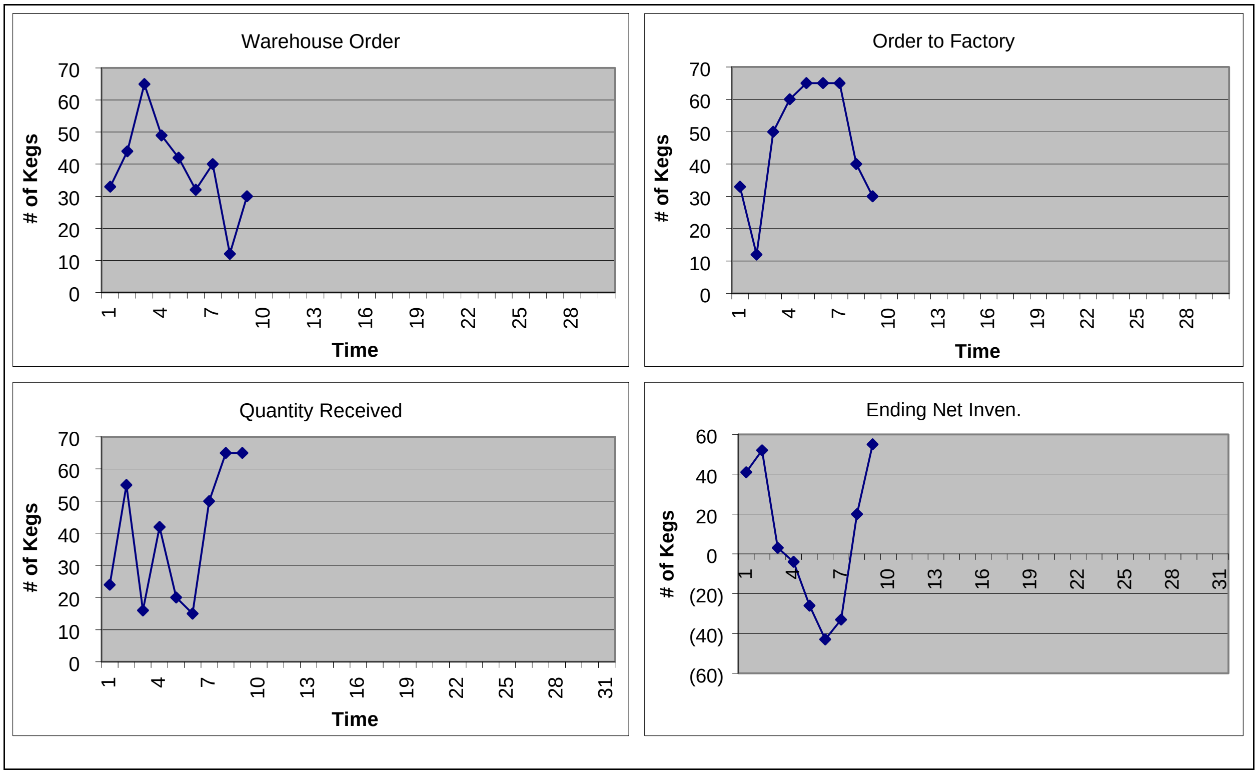

Figure 10. In-Game Graphs.

Players may choose to play a training version, which uses different demand values and moves at a slower pace. This allows them to learn at their own pace and practice any strategies they may have devised before playing "the real thing."

Benefits to Instructors:

From the instructor's point of view, setting up the game for use is simple. While the game comes "ready to play" as is, the instructor is free to change many of the game's parameters including demand and cost parameters and the game length. Once the parameters are set, the instructor need only distribute the program to each player. Password protection ensures that the parameters won't be tampered with. No other materials are needed during play. This eliminates the need for constructing game boards, procuring stacks of index cards and gathering mounds of change needed in the original beer game.

Since the program guides players through the steps of the game and automates the bookkeeping tasks, the need for instructor supervision is minimal. In fact, once players have been introduced to the game and shown how to use it, it is possible for them to play the game on their own, outside of class. Those familiar with the original beer game can appreciate the reduced handholding requirements.

At the end of a game, the program records the results for each team onto a floppy disk, automatically generating results ready for classroom use. This allows for rapid turnaround of results.

It is our intention that *Supply Chain Simulations* be used as a virtual laboratory. The structure of the games and the customizability of the parameters allow many different hypotheses to be tested under controlled conditions. The ease with which data is recorded and compiled lets instructors build a comprehensive database of experimental results after only a few semesters.

References

Barilla SpA (A), Harvard Business School Case, N9-694-046. 1984.

Chen, F. (1996), "Decentralized Supply Chains Subject to Information Delays," *Management Science* (forthcoming).

Sterman, J. (1984), "Instructions for Running the Beer Distribution Game (D-3679)," Sloan School of Management, MIT.

Sterman, J. (1989), "Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment," *Management Science,* 35, 3, 321-339.

Van Ryzin, G. (1998), "Analyzing Inventory Cost and Service in Supply Chains," Teaching Notes, Columbia Business School.