Value of Information in Capacitated Supply Chains
Author(s): Srinagesh Gavirneni, Roman Kapuscinski, Sridhar Tayur
Reviewed work(s):
Source: *Management Science*, Vol. 45, No. 1 (Jan., 1999), pp. 16-24
Published by: INFORMS
Stable URL: http://www.jstor.org/stable/2634919
Accessed: 01/11/2011 02:28

Value of Information in Capacitated Supply Chains

Srinagesh Gavirneni • Roman Kapuscinski • Sridhar Tayur
Austin Product Center—Research, Schlumberger, Austin, Texas 78726
University of Michigan Business School, Ann Arbor, Michigan 48109
Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213

We incorporate information flow between a supplier and a retailer in a two-echelon model that captures the capacitated setting of a typical supply chain. We consider three situations: (1) a traditional model where there is no information to the supplier prior to a demand to him except for past data; (2) the supplier knows the (s, S) policy used by the retailer as well as the end-item demand distribution; and (3) the supplier has full information about the state of the retailer. Order up-to policies continue to be optimal for models with information flow for the finite horizon, the infinite horizon discounted and the infinite horizon average cost cases. Study of these three models enables us to understand the relationships between capacity, inventory, and information at the supplier level, as well as how they are affected by the retailer's $(S - s)$ values and end-item demand distribution. We estimate the savings at the supplier due to information flow and study when information is most beneficial.

(*Information Sharing; (s, S) Policy; Optimal Policies; Capacitated Production-Inventory Model; Infinitesimal Perturbation Analysis*)

1. Introduction

The industrial supplier-retailer relations have undergone radical changes in recent years as the philosophy behind managing manufacturing systems continues to be influenced by several Japanese manufacturing practices. As more organizations realize that successful in-house implementation of Just-in-Time alone will have limited effect, they are encouraging other members of their supply chain to change their operations. This has resulted in a certain level of co-operation, mainly in the areas of *supply contracts* and *information sharing*, that was lacking before. This is especially true when dealing with customized products, and is most commonly seen between suppliers and their larger customers (retailers).

Our primary motivation to develop the models in this paper came from two sources. One was from the food industry, when a Pittsburgh based manufacturer—the largest supplier on the East Coast to a large fast-food chain—was asked to take the lead in integrated supply chain management. One of the goals of the manufacturer was to improve the understanding of the interactions between information, inventory, and capacity (because it affects customer service and cost), providing insights for the entire supply chain, consisting of 160 suppliers. Consequently, we take the supplier's point of view in this paper. The second (and a broader) motivation to study the benefits of information arose because of differing reactions to Electronic Data Interchange (EDI) benefits from industrial sources: While some were very happy with improved information, others were disappointed with its benefits (see Armistead and Mapes (1993) and Takac (1992)). Thus, while information is always beneficial, we wished to investigate when it is most beneficial and when it is only marginally useful. In the latter

0025-1909/99/4501/0016$05.00

case, we wanted to explore how improving some other characteristic of the system, such as end-item demand variance or supplier capacity might allow one to expect significant benefits from information.

The degree of cooperation varies significantly from one supply chain to another. The type of information sharing could range from generic (e.g. type of inventory control policy being used, type of production scheduling rules being used) to specific (e.g. day-to-day inventory levels, exact production schedules). We develop and analyze new models of these recent developments in information sharing.

Previous research in this area, i.e. incorporating information flow into inventory control and supply chains, has mainly followed two approaches. The first approach used by Scarf (1959), Iglehart (1964), Azoury (1985), Lovejoy (1990), and Lariviere and Porteus (1995, 1996) was to use the history of the demand process to more accurately forecast the demand distribution using Bayesian updates. The second approach, used in this paper, incorporates the information flow by developing new analytical models. This approach is used by Zheng and Zipkin (1990) to analyze the value of information flow in a two-product setting. They showed that using the information about the outstanding orders of both products results in improvements of system performance. Zipkin (1995) extended this to a multiitem production facility. Hariharan and Zipkin (1995) incorporated information about order arrivals, termed demand lead-times. Chen (1995) studied the relative benefits of echelon-stock policies over those of installation-stock policies in a multi-echelon environment. The latter decisions at a given facility depend only on local inventory information as opposed to this information combined with information on other facilities. A different way of expressing information is through relations among demand states. A stream of papers (e.g. Song and Zipkin (1993) and Sethi and Cheng (1997) models the demand process as dependent on exogenous ("informational") variables which themselves fluctuate according to a given stochastic process. Traditional capacitated models have recently received attention: Federgruen and Zipkin (1986a, b) provide the optimal policy for the stationary demand case; Glasserman and Tayur (1994, 1995) study multiechelon systems operated via a base stock policy; Aviv and Federgruen (1997), Kapuscinski and Tayur (1998) and Scheller-Wolf and Tayur (1998) study nonstationary demand settings.

In this paper, we study partial and complete information sharing in a supplier-retailer setting, and also compare these to a base case of no information. The retailer faces positive integer demands which are i.i.d. in any period ("the end-item demand") and places orders with the supplier according to an (s, S) policy (i.e., at the supplier-retailer interface there is an implicit fixed ordering cost). In a traditional setting (Model 0), the supplier is unaware of the retailer's demand or ordering policy, and merely observes the orders from her, i.e. he assumes that the demands from the retailer are i.i.d. in each period (we illustrate them as originating from a black box, cf. Figure 1). In Model 1, the supplier knows the demand distribution faced by the retailer, the fact that orders are placed according to an (s, S) policy, and the specific parameters s and S (see Figure 2). Finally, in Model 2 the supplier receives, in addition, immediate (periodic) information about the retailer's demands, perhaps via EDI links (see Figure 3).

Some of our qualitative results are as follows. It is intuitive that the optimal policy in a model with additional information should perform better than the optimal policy in a model with restricted information, and this is confirmed by our experimental results. However, when the end-item demand variance is high, or the value of $\Delta = S - s$ is very high or very

Figure 1 **Model 0—The Traditional Supply Chain**

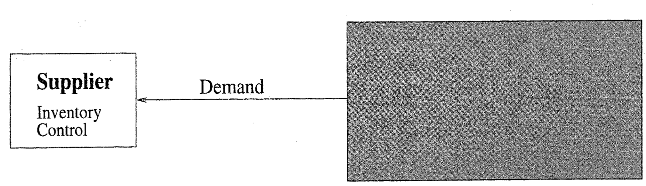

Figure 2 **Model 1—Supply Chain with Some Information Flow**

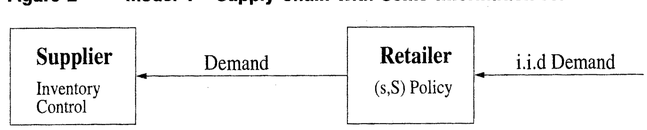

Figure 3 Model 2—Supply Chain with Full Information Flow

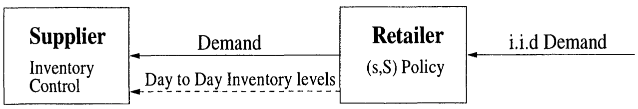

low, information is not that beneficial. On the other hand, if the end-item demand variance is moderate, and the value of $S - s$ is not extreme, the benefits of information are great. Similarly, information is not as beneficial if the supplier's capacity is low as compared to when it is high. With regard to penalty costs (at the supplier), the benefits of information appear to first increase and then drop off. Thus, before embarking on an EDI implementation, the supplier should verify his capacity and the fixed order cost faced by the retailer (as reflected by the retailer's Δ), while the retailer should attempt to reduce the end-item demand variance. When Δ and the demand variance are reduced to moderate values, and the supplier has moderate to high capacity, the information link is expected to be most beneficial. We justify these insights and other results in subsequent sections.

The rest of the paper is organized as follows: In §2 we describe the models and present structural properties of the optimal order up-to levels. Section 3 contains the computational results and insights. As a by-product, we obtain a simple heuristic for computing inventory levels. We conclude in §4 by summarizing our findings and discussing future directions of research.

2. Models

We consider a periodic review inventory control problem. The sequence of events in every period is as follows. First, the supplier decides on his production quantity for the period. Next, the retailer realizes her demand for the period. After satisfying (fully or partially) the demand, if her ending inventory level is below s, she places an order with the supplier to bring her inventory level to S. This order arrives at the beginning of the next period. If the supplier cannot satisfy the full order of the retailer, we assume that the retailer acquires the missing part of the order elsewhere. All this happens with no lead time. The (s, S) policy is optimal for the retailer, for example, when she incurs fixed plus linear ordering costs, linear holding and backlogging costs (under full backlogging) and if she faces i.i.d. demands, see e.g. Scarf (1962).

The supplier incurs linear holding and penalty costs (for portions of demand not satisfied from inventory), at unit rates h and b respectively. Every unit purchased costs a given, constant amount, which therefore has no impact on the optimal policy under the long run average cost criterion.

2.1. Model 0

In Model 0, the supplier has no information about the retailer, except what is available from past demand data. There are a number of ways to use this demand data to understand the demand process. Indeed, if the supplier senses that a (s, S) policy is being used, estimation of its parameters from data could lead to a model very close to Model 1 (described below) without obtaining any additional information from the retailer. However, for the purpose of studying the value of information, we assume that the supplier follows a naive approach and assumes that the demands from the retailer are following an i.i.d. process (that is, assuming a "traditional" model). Under this assumption, it is optimal for the supplier to use a modified order up-to policy with the same order up-to level in every period.[1]

We compute the appropriate order up-to level using infinitesimal perturbation analysis (IPA). By using IPA, we avoid having to make any assumptions on the distribution of the demands and at the same time account for the presence of finite capacity. The i.i.d. assumption considerably simplifies the model, and can be analyzed as in Glasserman and Tayur (1994, 1995).

2.2. Model 1

At the beginning of each period, the supplier knows the number of periods i that have elapsed since the last order was placed. Recall that we have assumed the end-item demand is at least one unit per period;

[1] A different straw man model is studied in Gavirneni (1997), and the results obtained there are qualitatively similar to those described in this paper.

therefore, the states are $i = 0, 1, \ldots, \Delta - 1$. Given this "external" state variable, future demands at the supplier are independent of past demands. Because the retailer obtains all of the ordered parts (some possibly from other sources), she brings her inventory position to S (after ordering) independent of the supplier's inventory position. Since the supplier knows that the retailer follows an (s, S) policy and since he knows the value of $\Delta = S - s$, he can determine the probability p_i that an order will be generated at the end of the coming period and the *cdf* $\Phi_i(\cdot)$ of the order size, given that an order is placed. Note that (with d_i end-item demand in period i)

$$p_i = \text{Prob}\{d_1 + d_2 + \cdots + d_{i+1} \geq \Delta | d_1 + \cdots + d_i < \Delta\}, \quad \text{and}$$

$$\Phi_i(x) = \text{Prob}\{d_1 + d_2 + \cdots + d_{i+1} \leq x | d_1 + \cdots + d_i < \Delta \quad \text{and} \quad d_1 + d_2 + \cdots + d_{i+1} \geq \Delta\}.$$

We further assume that $p_i \leq p_{i+1}$ for all $i < \Delta - 1$, $p_{\Delta-1} = 1$, and that $\Phi_i(\cdot) \geq \Phi_{i+1}(\cdot)$, i.e., the random variable associated with $\Phi_i(\cdot)$ is stochastically smaller than that associated with $\Phi_{i+1}(\cdot)$. This monotonicity assumption is satisfied if the end-item demand to the retailer has an Increasing Failure Rate (IFR) distribution (such as Uniform, Normal or Erlang distributions); see Gavirneni (1997) for a proof.

We briefly note some results for Model 1. As defined in Federgruen and Zipkin (1986a, b), a modified order up-to policy with level z is one where: If the inventory level is less than z, we raise it to z; if this level cannot be reached, we exhaust the available capacity; if the inventory level is above z, we produce nothing.

PROPERTY 1. *For finite and infinite horizons (discounted or average costs) an order up-to policy is optimal.*

PROOF. The proof follows standard steps as in Federgruen and Zipkin (1986a, b) and Kapuscinski and Tayur (1998). See Gavirneni (1997). □

In the infinite horizon the optimal order up-to level in a period depends only on the state of the system in that period. Let z_i be the infinite horizon optimal order up-to level in state i. In the finite horizon, the optimal order up-to level also depends on the remaining horizon length. Let y_i^n be the optimal order up-to level in state i, when n periods are left until the end of the horizon.

PROPERTY 2. *Optimal order up-to levels are increasing in i for both finite horizon and infinite horizon (discounted or average cost) models, i.e.,* $y_i^n \leq y_{i+1}^n$ *and* $z_i \leq z_{i+1}$.

PROOF. The proof follows from the monotonicity of p_i and $\Phi_i(\cdot)$. For details see Gavirneni (1997). □

We compute the optimal order up-to levels via simulation based optimization (using infinitesimal perturbation analysis, IPA). As in Glasserman and Tayur (1994) and Kapuscinski and Tayur (1998) we write the simulation recursions and differentiate them. Validation of the procedure follows standard steps for both the finite and infinite horizon. For the infinite horizon, we exploit the monotonicity of order up-to levels to show that regeneration occurs if and only if

$$z_{i+1} \leq \max\{z_i + C, [z_{\Delta-1} - \Delta]^+ + C\}$$

for $i < \Delta - 1$ where $z_0 = 0$ and $x^+ = \max(0, x)$. See Gavirneni (1997).

2.3. Model 2

At the beginning of each period, the supplier knows j = the number of units sold by the retailer since her last order ($j = 0, 1, \ldots, \Delta - 1$). Clearly, if j is known, having knowledge of the number of periods since the last order is of no incremental value. Moreover, given the "external" state variable j, future demands for the supplier are again independent of past demands. The supplier can determine p_j, the probability that an order will be placed at the end of the coming period, given the period starts in state j, as well as $\Phi_j(\cdot)$, the cdf of order size, given an order is placed:

$$p_j = \text{Prob}\{d \geq \Delta - j\}$$

and

$$\Phi_j(x) = \text{Prob}\{d \leq x - j | d \geq \Delta - j\}.$$

It is easily verified that $p_j \leq p_{j+1}$, for all $j < \Delta - 1$, $p_{\Delta-1} = 1$, and $\Phi_j(\cdot) \geq \Phi_{j+1}(\cdot)$, i.e., the probability

Figure 4 The Experimental Setup

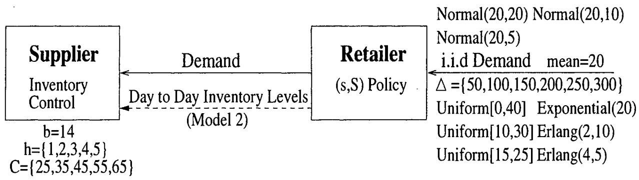

and quantity of realization of demand increases with j. Here we do not need the IFR assumption on the end-item demand distribution for monotonicity.

All the properties of Model 1 also hold for Model 2. Briefly: (1) the cost function is convex and the optimal policy is modified order up-to; (2) the order up-to levels are ordered; and (3) IPA can be used to find the optimal order up-to levels.

Figure 5 Plot of % Benefit Versus Capacity Between Model 0 and Model 1

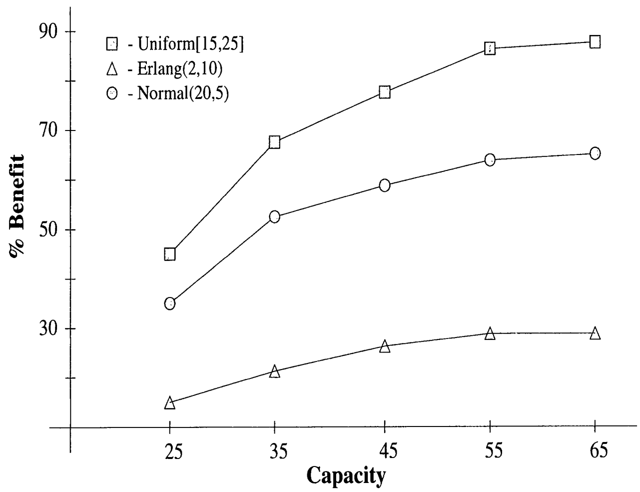

3. Computational Results

In this section we implement the solution procedures developed previously for all three models. Our goal is to understand the trade-offs between inventories, capacities, and information. We vary h, Δ, and the variance of end-item demand. (Since proportionally increasing both b and h does not change the results, it is sufficient to vary only one of them.)

The experimental design is shown in Figure 4. The values of b and mean demand were held constant at 14 and 20 respectively. The holding cost h was varied from 1 to 5 in increments of 1, while the capacity was varied from 25 to 65 in increments of 10. The value of Δ varied from 50 to 300 in increments of 50. We used nine demand distributions (all with mean 20 and various standard deviations σ): Uniform[0, 40] ($\sigma = 11.55$), Uniform[10, 30] ($\sigma = 5.77$), and Uniform[15, 25] ($\sigma = 2.87$); Exponential(20) ($\sigma = 20$), Erlang(2, 10) ($\sigma = 14.14$), and Erlang(4, 5) ($\sigma = 10$); Normal(20, 13.5) ($\sigma = 13.5$), Normal(20, 9.5) ($\sigma = 9.5$), and Normal(20, 5) ($\sigma = 5$). (Note: An Erlang(k, μ) results from adding k independent exponential distributions, each with mean μ, and a Normal(μ, σ) has a mean μ and standard deviation σ.[2]) There were a total of 675 instances in this computational setup.

3.1. Costs and Savings

We observed that for all the models the total cost increases if the holding cost or variance increases, or capacity decreases. These results are expected and so we will not elaborate on them further. We will look more closely at the savings realized due to additional information flow and their dependency on capacity, holding cost, variance, and Δ. The cost for Model 2 was always smaller than that for Model 1 which in turn was smaller than that for Model 0. This leads us to conclude that *information is always beneficial.*

First, we study the percentage savings realized between Models 0 and 1: the benefits of *some* information. The graph of these savings versus capacity is given in Figure 5 for a representative sample of demand distributions. (See Gavirneni (1997) for graphs relating percentage benefits to holding cost, Δ, and end-item demand variance.) The savings vary from 10% to 90%, with an average around 50%; they increase with capacity. These savings are largely due to the simplistic nature of Model 0. The stationary modified order up-to policy, optimal for Model 0, requires inventories to be held in every period. Model

[2] While generating demands using the Normal distribution, the nonpositive demands were discarded. To account for this, we selected the non-truncated mean so that the accepted demands had a mean of 20. The standard deviations of the observed demands for these three cases were 13.5, 9.5, and 5 and these values were used while analyzing the results.

Figure 6 Plot of % Benefit Comparing Model 1 to Model 2 Versus Capacity

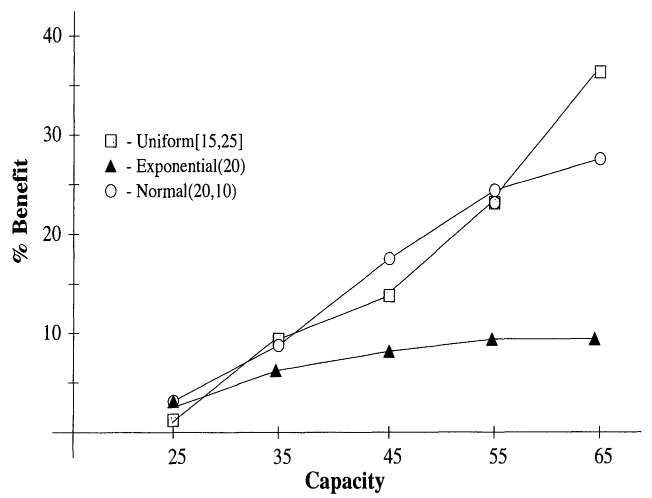

Figure 7 Plot of % Benefit Versus the Ratio of Penalty to Holding Cost Between Model 1 to Model 2

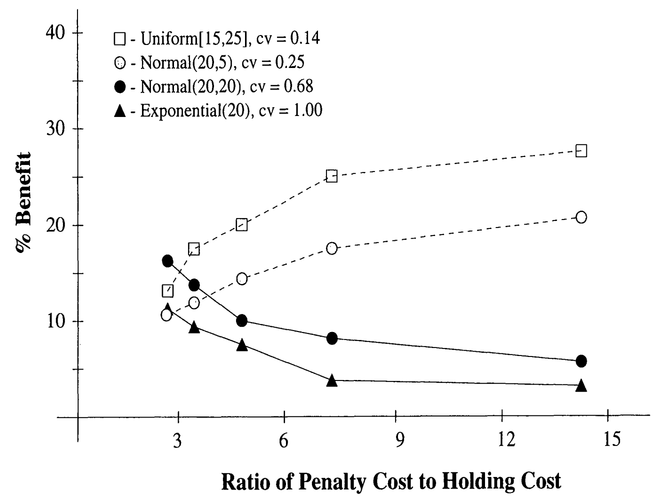

Figure 8 Plot of % Saving from Model 1 to Model 2 Versus Standard Deviation

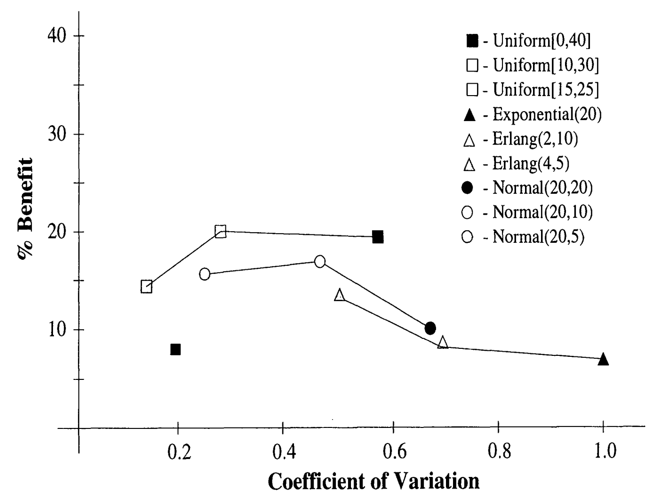

1 results in lower costs, in particular at higher capacities, since information in state i is used to schedule production more efficiently. However, reasonable benefits are noticed even at low capacities. We also observed that the benefit increases with the holding cost rate; see Gavirneni (1997) for details.

We next study the percentage savings realized between Model 1 and Model 2: the benefits of increased (and *full*) information. The plots of these savings with respect to capacity, ratio of penalty to holding cost, and standard deviation are given in Figures 6, 7, and 8 respectively. These savings range widely, from 1% to 35%. Details are described below.

1. *Effect of Capacity.* For every value of capacity and each demand distribution we averaged the percentage savings over all the holding costs. A representative sample of these results is presented in Figure 6. It is easily noticed that the percentage savings increase with capacity, and in contrast to what we noticed in Figure 5, there is almost no benefit at low capacities: When the capacity is 25, since the end-item mean demand is 20, one does not have much choice other than to produce in every period (the $S - s$ value does not really affect this) to meet the demand. However, when the capacity is high, the supplier has flexibility and thus can use the information to delay production or produce a larger quantity in a given period (if necessary).

2. *Effect of Ratio of Penalty to Holding Cost.* First, let us look at extreme values of the ratio. If the ratio is very close to 0, the supplier will not produce before seeing the demand, while a huge ratio makes it beneficial to keep extremely high inventory all the time due to finite capacity faced by the supplier. In both cases, additional information would have very limited benefit. Figure 7 shows the information benefits for moderate ratios. Interestingly, the range in which we observe increase (or decrease) in the benefit

of information appears to be highly dependent on the coefficient of variation: For small coefficients of variation, the benefit of information is observed to increase for a broad range of ratios (and a peak is reached at high penalty cost), while for high coefficients of variation, a peak is reached at a small penalty.

3. *Effect of End-Item Demand Distribution and Standard Deviation.* First, note that for standard deviation $\sigma = 0$, there is no difference between Models 1 and 2, and, therefore, % Benefit = 0. Figure 8 contains the plot of percentage savings versus standard deviation for each of the demand distributions. For the Uniform distribution, as the coefficient of variation was decreased from 0.578 to 0.288 to 0.144 the percentage savings first increased and then decreased. Similarly, for the Normal distribution as the standard deviation was changed from 0.675 to 0.475 to 0.25 the percentage savings initially increased from 7.55% to 9.37% and then dropped to 7.37%. For the Erlang distributions, as the coefficient of variation was decreased from 1 to 0.707 to 0.5 the percentage savings increased. It appears that as standard deviation increases, the percentage benefit drops. This could be due to the fact that, as in item 2 above, when the variance is very high the reduction in uncertainty due to additional information at the supplier is not significant. In other words, the information available relative to the overall system uncertainty is small, and thus does not reduce costs as effectively. Thus, we observe that *information is most beneficial at moderate values of variance.*

4. *Effect of* Δ. The percentage savings between Models 1 and 2 as a function of Δ are displayed in Figure 9 for two distributions, Exponential(20) and Erlang(4, 5). For both distributions, as the value of Δ is increased these savings increase initially and then start to decrease. This behavior can be explained as follows. When Δ is large (compared to capacity and mean end-item demand), due to finite capacity, in anticipation of a large order we must start building up inventory over time, effectively reducing the benefit of information. On the other hand, when Δ is small (relative to mean end-item demand), the end-item demand is passed through to the supplier almost every period and there is no significant difference between Models 1 and 2. So it appears that *information is less beneficial at extreme values of* Δ.

Figure 9 **Plot of % Savings from Model 1 to Model 2 Versus Δ**

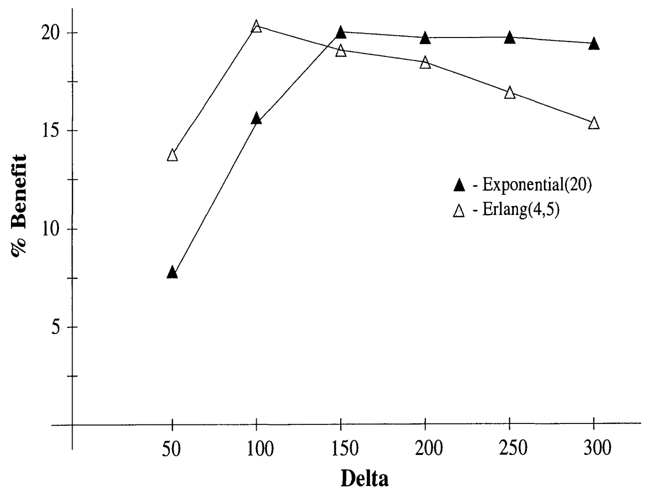

Figure 10 **The Optimal Order Up-to Levels for exp(20), $h = 5$, Model 1**

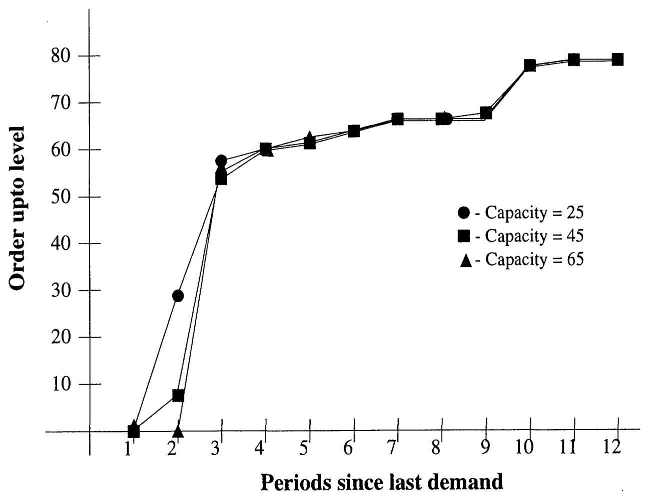

3.2. Inventory Levels

The optimal order up-to levels for Model 1 for the case of Exponential(20), $h = 5$ are given in Figure 10. When $C = 65$, the order up-to levels are zero for the two periods immediately following a demand. In period 3 the order up-to level jumps to 55 and from then on slowly climbs to around 80. In fact we can show that:

PROPERTY 3. *When* $C = \infty$ *the optimal order up-to levels satisfy*

1. $z_i = 0$, *for all* $i < i^*$;
2. $z_{i^*} \geq \Delta$; *and*
3. $z_k \leq \Delta + \gamma$,

where $i^* = \min\{i | p_i > h/(h + b)\}$, $\Psi(\cdot)$ *is the end-item demand distribution, and* $\gamma = \inf\{y | \Psi(y) \geq b/(h + b)\}$.

PROOF. See Gavirneni (1997). □

Next observe the optimal order up-to levels for other values of capacity: The optimal order up-to levels do not change significantly in periods 3 and above. The major changes occur, if at all, in periods immediately after the demand to facilitate reaching the required inventory level for period 3 ($= i^*$). Similar observations regarding the order up-to levels apply to other distributions. Using these observations we can draw the following insight into how one can approximate order up-to levels in the finite capacity case without IPA: Solve a given finite capacity system as if it had infinite capacity (we provide a quick recursive procedure to compute the optimal order up-to levels for infinite capacity in Gavirneni (1997) to obtain the order up-to levels $\{z_i\}$ particularly i^* and z_{i^*}. To obtain the order up-to levels $\{z_i^c\}$ for finite capacity, set $z_i^c = z_i$, $i \geq i^*$ and $z_{i-1}^c = (z_i^c - C)^+$, for all $i < i^*$, so that the order up-to level in i^* can be reached under the capacity restriction.

4. Conclusions

In this paper we have incorporated information flow into inventory control models. This gave rise to interesting nonstationary demand processes. Optimal policy structures are order up-to policies. An extensive computational study provided us with insights on the savings and relative benefits due to information flow. These insights were useful in the design of a major fast-food company's (FFC) supply chain, and we outline the setting briefly.

Due to extraordinary performance (in terms of on-time delivery and quality of products) by a food manufacturer (NSP) for two consecutive years, FFC asked NSP to manage the other 160 or so suppliers. Thus, FFC "outsourced" supplier management to NSP. A new organization called H1 Logistics was created. Among its many tasks was the following: to determine whether EDI links between FFC warehouses and suppliers would be beneficial.

As part of this effort, based on transportation times, type of product, end-item demand, volumes, and capacity, we first decided which supplier would ship directly to a FFC warehouse, and which supplier would ship it to NSP warehouse (which would, after consolidation, ship to a FFC warehouse). Similarly, the period length was determined, and this was either one day, one week or 2 weeks. Overnight transportation is possible between some suppliers and NSP, between NSP and some FFC warehouses, and between some suppliers and some FFC warehouses. The longest transportation time between locations in the supply chain is under 2 days, typically over the weekend. Due to judicious selection of period length and choice of transportation, our periodic review model with zero lead time was considered adequate. Furthermore, FFC expects a near perfect supply, and suppliers either work overtime or shift capacity pre-allocated to other brands to meet this requirement. This additional cost, that is supplier dependent, is captured as penalty cost. The holding costs are also different across suppliers. Based on relative values of holding and penalty costs, the suppliers were classified as low, medium or high in terms of holding and penalty costs.

Similarly, based on capacities of suppliers relative to the period volumes (end-item demand at the warehouse, for example), and the high fixed costs of ordering, we classified suppliers as having low (high) capacity and warehouses as having a low (high) Δ value. This helped us understand when information would be most useful, and also how we could further improve the supply chain performance. While several real-world concerns such as correlation and production variabilities were ignored in our models, the management at NSP and H1 Logistics viewed the insights gained as useful in their decision making; for example, EDI links were added between FFC warehouses and those suppliers with medium to high capacity and for products with moderate demand variation.

Future work is planned in four directions: (1) performing similar studies of two stage models under a less restrictive setting; (2) extending this analysis to multiple retailers; (3) studying inventory-information-capacity interactions in a supply chain with more than two stages; and (4) studying, at a theoretical level, more general demand processes[3] at the supplier, where the transition from one state to another follows a general probability matrix, of which the current non-stationary situations are special cases.[4]

[3] We thank Professor George Shanthikumar for this suggestion.

[4] We have incorporated many of the suggestions by the Departmental Editor and the Associate Editor that have improved the presentation of this paper. We also thank the referees for their comments and suggestions. Alan Scheller-Wolf's editorial help has been most invaluable.

References

Armistead, C. G., J. Mapes. 1993. The impact of supply chain integration on operating performance. *Logistics Inform. Management* **6**(4) 9–14.

Aviv, Y., A. Federgruen. 1997. Stochastic inventory models with limited production capacity and periodically varying parameters. *Probab. Engrg. Inform. Sci.* **11**(2) 107–135.

Azoury, K. S. 1985. Bayes solutions to dynamic inventory models under unknown demand distributions. *Management Sci.* **31**(9) 1150–1160.

Chen, F. 1995. Echelon reorder points, installation reorder points, and the value of centralized demand information. Working Paper Graduate School of Business, Columbia University, New York.

Federgruen, A., P. Zipkin. 1986a. An inventory model with limited production capacity and uncertain demands I: The average-cost criterion, *Math. Oper. Res.* **11**(2) 193–207.

——, ——. 1986b. An inventory model with limited production capacity and uncertain demands II: The discounted-cost criterion. *Math. Oper. Res.* **11**(2) 208–215.

Gavirneni, S. 1997. Inventories in supply chains under cooperation. Ph.D. Thesis, Carnegie Mellon University, Pittsburgh, PA, September.

Glasserman, P., S. Tayur. 1994. The stability of a capacitated, multi-echelon production-inventory system under a base-stock policy. *Oper. Res.* **42**(5) 913–925.

——, ——. 1995. Sensitivity analysis for base-stock levels in multi-echelon production-inventory systems. *Management Sci.* **42**(5) 263–281.

Hariharan, R., P. Zipkin. 1995. Customer-order information, lead times, and inventories. *Management Sci.* **41**(10) 1599–1607.

Iglehart, D. L. 1964. The dynamic inventory problem with unknown demand distributions. *Management Sci.* **10** 429–440.

Kapuscinski, R., S. Tayur. 1998 A capacitated production-inventory model with periodic demand. *Oper. Res.* **46**(6) 899–911.

Karlin, S. 1968 *Total Positivity Vol. I*. Stanford University Press, Stanford, CA.

Lariviere, M.A., E. L. Porteus. 1995. Manufacturer-retailer contracting under an unknown demand distribution. Graduate School of Business Working Paper, Stanford University, Stanford, CA, December.

——, ——. 1996. Stalking information: Bayesian inventory management with unobserved lost sales. Graduate School of Business Working Paper, Stanford University, Stanford, CA, December.

Lovejoy, W.S. 1990. Myopic policies for some inventory models with uncertain demand distributions. *Management Sci.* **36**(6) 724–738.

Scarf, H. 1959. Bayes solutions to the statistical inventory problem. *Ann. Math. Statist.* **30** 490–508.

——. 1962. The Optimality of (s, S) policies in the dynamic inventory problem. K.J. Arrow, S. Karlin, P. Suppes, eds. *Mathematical Methods in Social Sciences*. Stanford University Press, Stanford, CA.

Scheller-Wolf, A., S. Tayur. 1998. A Markovian dual-source production-inventory model with order bands. GSIA working paper, CMU, Pittsburgh, PA.

Sethi, S., F. Cheng. 1997 Optimality of (s, S) policies in inventory models with Markovian demands. *Oper. Res.* **45**(6) 931–939.

Song, J., P. Zipkin. 1993. Inventory control in a fluctuating demand environment. *Oper. Res.* **41**(2) 351–370.

Takac, P.F. 1992. Electronic data interchange (EDI): an avenue to better performance and the improvement of trading relationships? *Internat. J. Computer Applications in Technology* **5**(1) 22–36.

Zheng, Y., P. Zipkin. 1990. A queuing model to analyze the value of centralized inventory information. *Oper. Res.* **38**(2) 296–307.

Zipkin, P. 1995. Performance analysis of a multi-item production-inventory system under alternative policies. *Management Sci.* **41**(4) 690–703.

Accepted by Awi Federgruen; received March 9, 1996. This paper has been with the authors 3 months for 3 revisions.