SIMULATION OF SCHEDULED ORDERING POLICIES IN DISTRIBUTION SUPPLY CHAINS

Lucy G. Chen

NUS Business School
National University of Singapore
Singapore 117592, SINGAPORE

Srinagesh Gavirneni

Johnson Graduate School of Management
Cornell University
Ithaca, N.Y. 14853, U.S.A.

ABSTRACT

In this paper we study a decentralized distribution supply chain with one supplier and many newsvendor-type retailers that face exogenous end-customer demands. Using total supply chain cost as our primary measure of performance, we compare two scheduled ordering policies – Balanced ordering and Synchronized ordering – with the traditional newsvendor-type ordering behavior. Via the use of simulation, we evaluate the effectiveness of the two scheduled ordering policies, and identify how the performance of the scheduled ordering policies changes with different supply chain parameters, such as the number of retailers, the supplier's expediting cost, the supplier's capacity limit, *etc*.

1 INTRODUCTION

Distribution supply chains, which consist of one supplier serving many retailers that face exogenous end-customer demand, are very hard to manage because of the inherent multiple sources of uncertainty. Uncoordinated retailer ordering behavior magnifies the complexity of these distribution systems. Limited capacity at the supplier further complicates the problem as the supplier needs to determine how to allocate the limited inventory among the retailers. Anupindi, Bassok and Zemel (2001) and Chen, Federgruen and Zheng (2001) both illustrated the complexity of analyzing decentralized distribution systems.

Different approaches have been tried to smooth the flow of materials and information in distribution systems. Strategies that restrict retailer behavior have become popular in practice. For example, GATX, a third party logistics company that serves 6500 BP and Chevron service stations, divides the service stations into 5 groups, each of which is required to order only on a preassigned day of the week. So 20% of the service stations order on Monday, another 20% order on Tuesday, etc. (Andel 1995). As another example, the company described in the case "Norton Auto Supply" (Hammond 1989) partitions the regional distribution centers into five groups, each of which is served on one separate day of a week. Strategies like this intuitively reduce the demand variability for the supplier. Yet there has been no rigorous research on the effectiveness of these restrictive strategies from the total supply chain perspective. In this paper, we try to fill this void via the use of simulation.

Scheduled ordering policies have been studied, albeit without any consideration of information sharing, by other researchers. By studying the demand variance that the supplier faces, Lee, Padmanabhan and Whang (1997) found that balanced ordering (same number of retailers order in every period) is preferable to randomized ordering (retailers may order in any random period), which outperforms synchronized ordering (all retailers order in the same period). Their research, however, did not evaluate the total supply chain performance, which is our main focus. By restricting the retailers to order only in batches and only in infrequent review periods, Cachon (1999) numerically found that switching from synchronized ordering to balanced ordering reduces the supply chain holding and backorder costs. His research, however, did not consider information sharing.

The problem setting in our paper is different from the ones in the previous literature. The retailers in our setting do not face fixed ordering cost. Therefore they may order freely in every period. Under our scheduled ordering policies we still allow shipments in those periods that retailers cannot order freely. Since the special case with $\delta = 0$ would capture the setting where no shipment is allowed in the non-ordering periods, our work generalizes existing research. Our setting is similar to that in Zhu, Gavirneni and Kapuscinski (2007) which, however, focused on a serial system with one supplier and one retailer.

The objective of this paper is two-fold. We want to tabulate the magnitude of the cost reduction that could be achieved by the scheduled ordering policies when compared to the traditional newsvendor-type ordering behavior. Second, we want to study how different supply chain parameters, such as demand variance, demand correlation level, number of retailers, capacity limit, and supplier's expediting and holding costs, would impact the performance of the scheduled ordering policies.

The rest of the paper is organized as follows. In Section 2, we introduce the basic problem setup and describe the ordering policies we are looking at. In Section 3, we examine the performance of two scheduled ordering policies and compare it with the traditional newsvendor-type ordering behavior. We also compare the performance of the two scheduled ordering policies. In Section 4, we conclude the paper.

2 PROBLEM DESCRIPTION

Consider a decentralized distribution supply chain with one supplier and m identical newsvendor-type retailers. The retailers have the same cost parameters and face identical end-customer demands in each period. The end-customer demands are assumed to be *iid over time* with cumulative distribution function (cdf) $\Psi(\cdot)$ and probability distribution function (pdf) $\psi(\cdot)$. So whenever we mention demand correlation, we mean *correlation across the retailers*. For ease of exposition, we will consistently call the demand that the supplier faces "retailer demand" to differentiate that from end-customer demand. The inventory holding and penalty costs at each retailer are h_r and π_r. The inventory holding cost at the supplier is h_s. The supplier incurs an expediting cost π_s, which will be made clear in following paragraphs, if she does not have enough inventory to satisfy retailer demands. There is no fixed ordering cost involved. Information sharing is considered. By information sharing, we mean the supplier knows the realized end-customer demand.

We evaluate this supply chain in a one-product periodic-review setting. With the total supply chain cost per period as our performance measure, we evaluate the following retailer ordering policies:

Free Ordering Policy (FOP) This is the traditional ordering policy where all the retailers can freely order whatever they want in any period. Retailers will generally order in every period, since they are typical newsvendor-type retailers with no fixed ordering costs.

Balanced Ordering Policy (BOP) This policy operates with an ordering cycle of m periods (m is equal to the number of retailers). In every period only one retailer can order freely. Each retailer may order freely only in one period of the cycle. In the remaining periods of the cycle, a predetermined fixed quantity δ is shipped to him.

Synchronized Ordering Policy (SOP) This policy operates with an ordering cycle of m periods (m is equal to the number of retailers). All retailers order freely in the same one period of the cycle. In the remaining periods of the cycle, a predetermined fixed quantity δ is shipped to every retailer.

It is clear that under BOP and SOP, retailers will incur a higher cost than they do under FOP, as they cannot follow the stationary order-up-to policy which is optimal for them.

Therefore, the benefit of the scheduled ordering policies, if any, could only come from the supplier's side. The key question is whether the benefit to the supplier is greater than the cost increase to the retailers. If so, then the supplier could transfer some of the benefit to the retailers and everyone would be better off.

The sequence of the events in every period is: (a) the supplier receives the units produced in the previous period; (b) retailers place orders if they can in that period; (c) if the supplier does not have enough inventory to fulfill all retailer demand, she gets the product from an outside source immediately at a higher expediting cost; (d) the supplier ships the product to every retailer, either at the quantity the retailer selects or δ; (e) the supplier decides how much to produce in this period; (f) the end-customer demands occur and retailers satisfy the demands as much as possible with their on-hand inventory. Unsatisfied demand at the retailers is backlogged; (g) inventory-related costs (holding costs for the supplier and retailers, backorder cost for retailers and expediting cost for the supplier) are tabulated.

In performing the analysis, we make the following assumptions: (i) The production at the supplier takes one period, which means one period of lead time. (ii) The supplier fulfills all retailer orders every period. If the supplier does not have enough inventory on hand, she will get the product from an outside source immediately at a higher expediting cost, and then ship it to the retailers right away. This implies a high service standard at the supplier side. (See Lee, So, and Tang 2000 for more discussion on the appropriateness of this assumption.) (iii) The retailers receive their shipments immediately.

3 SIMULATION RESULTS

Modified order-up-to policies are optimal for both the supplier and retailers. Due to the finite capacity restriction, the optimal order-up-to levels are generally not available in closed form (Federgruen and Zipkin 1986). However, they can be computed by using IPA (Kapuscinski and Tayur 1998).

3.1 Experimental Setup

To cover the cases when scheduled ordering policies are effective and when they are not, we test a wide range of supply chain parameter combinations. We fix $h_r = 1$ and $\pi_r = 9$. The experimental design on other supply chain parameters is presented in Table 1. We have ρ represent the end-customer demand correlation coefficient and ω capture the supplier's capacity limit. For example, $\omega = 1.5$ means the supplier's capacity is 1.5 times the mean end-customer demands for all the retailers in every period. We use erlang distributions for the end-customer demand. For example, $(2, 10)$ is the set of parameters for an erlang distribution with

mean 20 and variance 200. Notice that the erlang distribution with parameters $(1, 20)$ is equivalent to an exponential distribution with mean 20.

Table 1: Supply chain parameters for computational study.

h_s	{0.5, 0.75, 1, 1.25}
π_s	{9, 19, 39, 99}
N	{2, 3, 4, 5 }
$\mathscr{D}$	Erlang{(1,20), (2,10), (4,5), (8,2.5)}
ρ	{0, 0.25, 0.5, 0.75, 1.0}
ω	{1.5, 2.5, 3.5, 4.5, ∞}

We set the fixed shipment quantity δ equal to the mean end-customer demand. Notice that this choice of δ might not be an optimal solution. But Zhu, Gavirneni, and Kapuscinski (2007) found that it was a good approximation. We compute the percentage cost difference between the different ordering policies as following:

$$\Delta_{fb} = \frac{\Theta_f - \Theta_b}{\Theta_f} \times 100\%;$$
$$\Delta_{fs} = \frac{\Theta_f - \Theta_s}{\Theta_f} 100\%;$$
$$\Delta_{sb} = \frac{\Theta_s - \Theta_b}{\Theta_s} 100\%.$$

For example, a bigger Δ_{fb} indicates better performance of BOP compared to FOP. Given the different values of the supply chain parameters in Table 1, we tested $6,400$ different combinations for the case where there is information sharing. In reporting the results, we take the average over all the experiments related to a specific level of the parameter on which we are focused. For example, if we want to see the percentage cost difference when the number of retailers is 5, then we take an average of the percentage cost differences over all the $1,600$ experiments in which the number of retailers is 5.

3.2 Effectiveness of BOP

Let us first take a look at the overall performance of BOP. Figure 1 shows the histograms of the percentage cost reduction over all the experiments. 64% of the $6,400$ experiments have positive cost reduction, which could be as high as 72%, with an average of 16%. There is a large number of experiments under which BOP is not effective. We will see in the next section how different supply chain parameters impact the performance of BOP.

3.2.1 Effect of Supply Chain Parameters

In this section we will explore in detail how the performance of BOP is impacted by different levels of supply chain parameters.
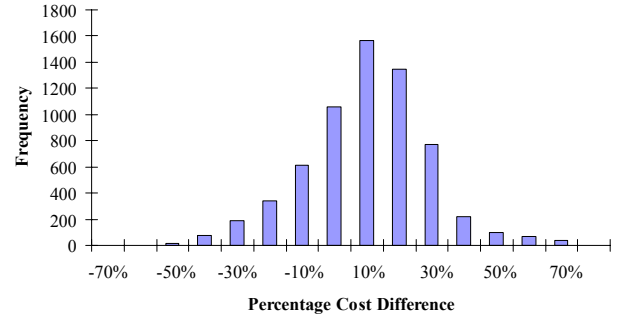

Figure 1: Histogram of the percentage cost reduction by using BOP.

Effect of Number of Retailers Figure 2 shows how the average percentage cost difference changes with the number of retailers. We observe that the average cost saving by using BOP could be as much as 23.22%. It appears that when the end-customer demand correlation level is low ($\rho = 0, 0.25$), BOP performs best when there are two retailers and it gets worse as the number of retailers increases. However, when the correlation level is relatively high ($\rho = 0.75, 1.0$), BOP performs best when there are three retailers. We conclude that *BOP is more effective when the number of retailers is small*.
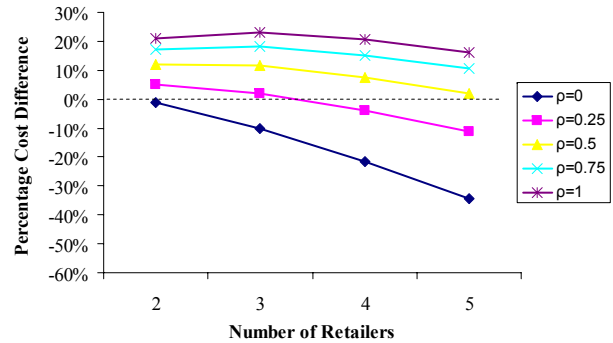

Figure 2: Percentage cost difference between BOP and FOP versus the number of retailers.

Effect of End-Customer Demand Variance The performance of BOP as a function of the end-customer demand variance is shown in Figure 3. The cost reduction by using BOP could be as much as 22.38% on average. The effect of demand variance is, however, not immediately obvious. It appears that when the end-customer demands are independent, the percentage cost difference tends to decrease as the variance increases. When the end-customer demands are correlated, the percentage cost difference increases as the variance increases. We thus conclude that *BOP is more effective with smaller demand variance when the demand is independent, whereas it is more beneficial with greater demand variance when the demands are correlated.*

Effect of the Supplier's Expediting Cost We illustrate the effect of the supplier's expediting cost in Figure 4. It is
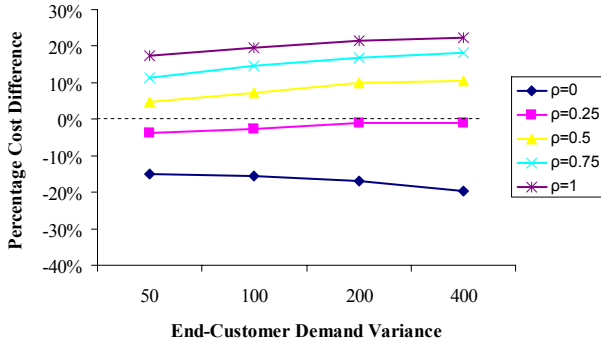
Figure 3: Percentage cost difference between BOP and FOP versus the end-customer demand variance.

observed that the average percentage cost reduction could be as much as 33.87%, and it increases as the supplier's expediting cost increases. The main reason is that as the supplier's expediting cost gets bigger, the cost saving realized by the supplier is getting higher. For example, with information sharing and fixed "$h_s = 1, \rho = 0.5, \omega = 1.5$ and Erlang $(2, 10)$," the supplier's cost reduction is 55.44% when $\pi_s = 9$ and 66.42% when $\pi_s = 99$. We thus conclude that *BOP is more effective when the supplier's expediting cost is higher*.
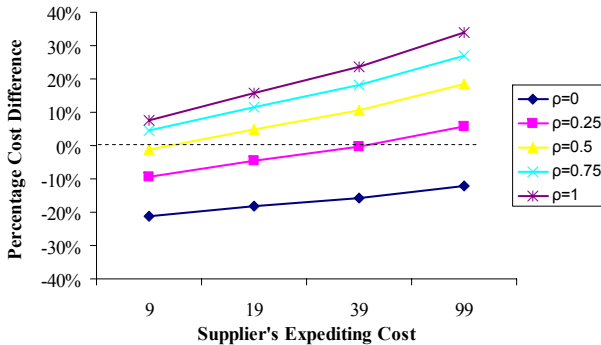

Figure 4: Percentage cost difference between BOP and FOP versus the supplier's expediting cost.

Effect of the Supplier's Holding Cost Figure 5 shows that the average percentage cost difference w.r.t. the supplier's holding cost could be as high as 26.87%. It appears that the percentage cost reduction increases as the supplier's holding cost increases. The reason for this behavior is similar to the one for the effect of the supplier's expediting cost. As the supplier's holding cost gets bigger, the cost saving realized by the supplier gets higher. For example, with information sharing and fixed "$\pi_s = 39, \rho = 0.5, \omega = 1.5$ and Erlang $(2, 10)$," the supplier's cost reduction is 57.38% when $h_s = 0.5$ and 59.26% when $h_s = 1.25$. We conclude that *BOP is more effective when the supplier's holding cost is higher*.

Effect of the Supplier's Capacity Limit Figure 6 illustrates the average cost difference as a function of the
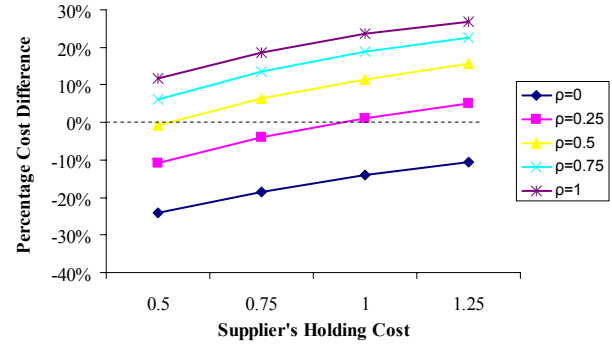

Figure 5: Percentage cost difference between BOP and FOP versus the supplier's holding cost.

supplier's capacity limit. Notice that the performance of BOP is better when the supplier has tighter capacity limit and could be as much as 32.11% on average. The reason for this behavior is that when the supplier has high capacity, she is able to cope with the demand variability better. Thus the advantage of using BOP is relatively diminished. However, the supplier's ability to react to the random demand is not that good if she has very tight capacity, which makes BOP significantly more beneficial. For example, if there is information sharing and we have fixed "$h_s = 1, \pi_s = 39, \rho = 0.5$ and Erlang $(2, 10)$," the supplier's cost reduction is 61.59% when $\omega = 1.5$ and 58.63% when $\omega = 4.5$. We conclude that *BOP is more effective when the supplier has tighter capacity*.
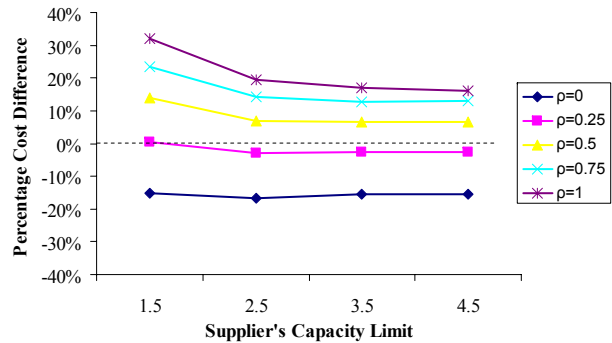

Figure 6: Percentage cost difference between BOP and FOP versus the supplier's capacity limit.

Effect of Demand Correlation From Figure 2 to 6, we have consistent observations that BOP performs better as the demand correlation level increases. The reason is as follows. The main benefit of BOP is reducing the variability faced by the supplier in every period. As the correlation among the end-customer demand increases, the variability under FOP compared to the one under BOP is getting larger. Hence the benefit of BOP is more significant when the demand correlation level is high. Here is an example. If there is information sharing and we have fixed "$h_s = 1, \pi_s = 39, \omega = 1.5$, and Erlang $(2, 10)$," the supplier's cost

reduction is 4.28% when $\rho = 0$ and 67.36% when $\rho = 1$. We hence conclude that *BOP is more effective when the demand correlation is higher*.

3.3 Effectiveness of SOP

SOP pools all demand risk into one period and leaves other periods risk free. This might be attractive under some scenarios, as on average the risk might be reduced. We observe that for the 6400 incidents with information sharing, SOP is effective for 1307 of them. Regardless of the small portion, there is a clear trend that the effectiveness of SOP decreases as the capacity limit gets tight. In fact, 74.13% (969) of the effective incidents are from the case when there is no capacity limit. The reason is probably as follows. SOP allows all customers to order freely in the same one period, which results in huge demand risk in one period. With limited capacity, the supplier will suffer from huge expediting cost, which cannot be compensated by the benefit she can possibly get from being risk-free in other periods. In view of this, we conclude that SOP is not effective when there is a tight capacity. We thus will only discuss the effectiveness of SOP for the case when there is no capacity limit.

Figure 7 shows the histogram of the percentage cost reduction over all the experiments of interest. Among the 1280 incidents with information sharing and no capacity limit, we find that SOP is effective 75.70% of the time. Compared to FOP, the cost reduction from using SOP can be as high as 41.2% with an average of 14.71%.
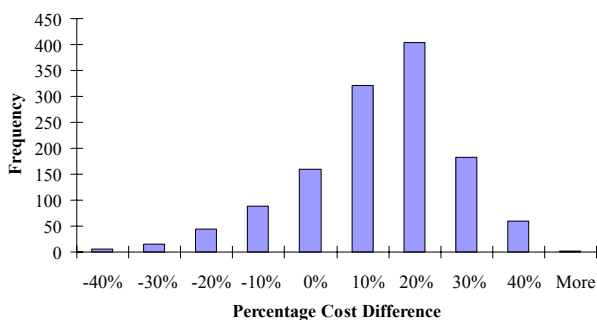

Figure 7: Histogram of the percentage cost reduction by using SOP.

Effect of Number of Retailers Figure 8 shows the performance of SOP with the number of retailers. We observe that the average cost saving by using SOP could be as much as 19.19%. Both the trend of the performance and the logic behind it are the same as the BOP case. We thus won't go into the details here.

Effect of End-Customer Demand Variance The performance of SOP as a function of the end-customer demand variance is shown in Figure 9. The average cost reduction by using SOP could be as much as 17.09%. The effect of
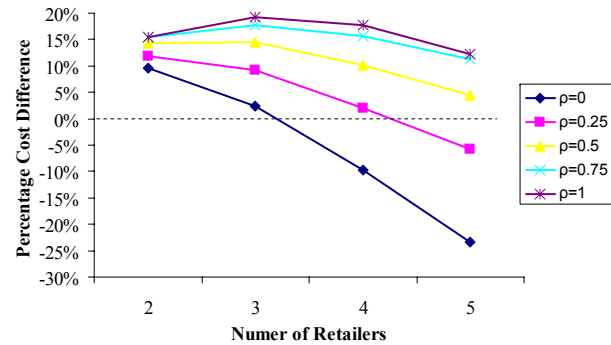

Figure 8: Percentage cost difference between SOP and FOP versus the number of retailers.

demand variance is, however, not immediately obvious. It appears that the percentage cost difference first increases with the end-customer demand variance, and then decreases. The reason is probably as follows. With the help of information sharing, when the demand variance is relatively small, the cost increase at the supplier is slower than the cost reduction at the retailer. Thus the performance of SOP first gets better with the increase of demand variance. However, when demand variance is very high, the information could not help much. The cost increase at the supplier will increase much faster than the cost reduction at the retailers. We thus conclude that *SOP is more effective with moderate demand variance*
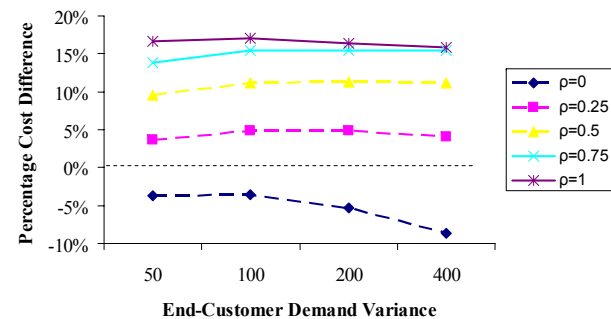

Figure 9: Percentage cost difference between SOP and FOP versus the end-customer demand variance.

Effect of the Supplier's Expediting Cost We illustrate the effect of the supplier's expediting cost in Figure 10. It is observed that the average percentage cost reduction could be as much as 24.83%. We observe that SOP performs better as the supplier's expediting cost increases. The logic behind this observation is the same as the BOP case. We thus conclude that *SOP is more effective when the supplier's expediting cost is higher*.

Effect of the Supplier's Holding Cost Figure 11 shows that the average percentage cost difference w.r.t. the supplier's holding cost could be as high as 24.72%. It appears that the percentage cost reduction increases as the supplier's holding cost increases. The reason behind this the same as
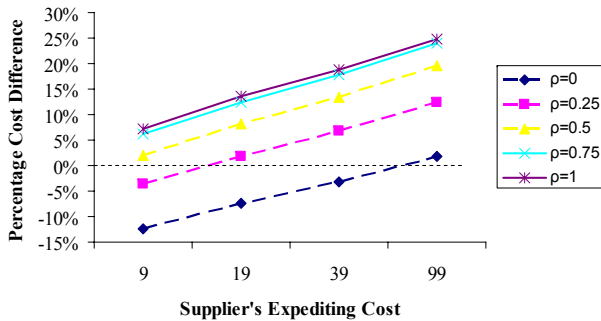
Figure 10: Percentage cost difference between SOP and FOP versus the supplier's expediting cost.
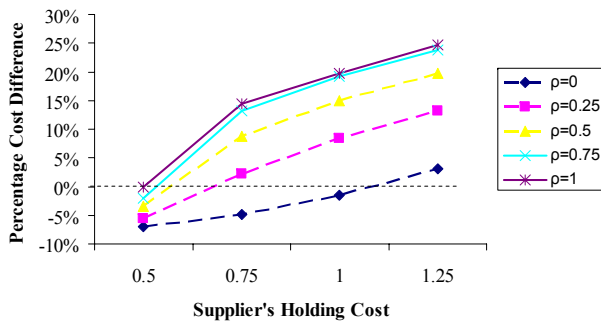

Figure 11: Percentage cost difference between SOP and FOP versus the supplier's holding cost.

the one for the BOP case. We conclude that *SOP is more effective when the supplier's holding cost is higher*.

Effect of the Demand Correlation The performance of SOP with the demand correlation has the same trend as that of BOP, that is, it performs better as the demand correlation increases. The reason for this behavior is also the same as the one for the BOP case.

3.4 BOP Versus SOP

In this section, we compare the effectiveness of BOP and SOP. Note that BOP is effective for many more incidents than SOP is. Since we conclude SOP is only worth looking at when there is no capacity limit and there is information sharing, we will compare it to BOP under this scenario. Among all the 969 incidents that SOP is effective, SOP outperforms BOP 84.93% of the time. The cost reduction can be as high as 20.79% with an average of 6.22%.

4 CONCLUSIONS

In this paper, via simulation, we studied the effect of two scheduled ordering policies (balanced ordering and synchronized ordering) in a decentralized distribution supply chain with one supplier and many newsvendor-type retailers. Information sharing was considered.

Through an extensive numerical study, we found that there was a large number of experiments in which BOP was not effective. However, there is a clear distinction between when BOP worked and when it did not. When it was effective, the cost reduction could be as much as 72%, with an average of 15%. For SOP, we found that it is only effective when there is no capacity limit, and that the cost reduction of using SOP instead of BOP can be as high as 20.79% with an average of 6.22%. We also observed that the scheduled ordering policies perform best when (i) the number of retailers is small; (ii) the supplier's expediting and holding costs are high; (iii) the end-customer demand variance and correlation level are high; and (v) when the supplier's capacity limit is tight. SOP shares most of the performance trend with BOP.

REFERENCES

Andel, T. 1995. There is power in numbers. *Transportation & Distribution* 36(8): 67-70.

Anupindi, R., Y. Bassok, and E. Zemel. 2001. A general framework for the study of decentralized distribution systems. *Manufacturing & Service Operations Management* 3(4): 349-368.

Cachon, G. 1999. Managing supply chain demand variability with scheduled ordering policies. *Management Science* 45(6): 843-856.

Chen, F., A. Federgruen, and Y. Zheng. 2001. Coordination mechanisms for a distribution system with one supplier and multiple retailers. *Management Science* 47(5): 693-708.

Federgruen, A. and P. Zipkin. 1986. An inventory model with limited production capacity and uncertain demands I: the average cost criterion. *Mathematics of Operations Research* 11(2): 193-207.

Gavirneni, S. 2002. Information flows in capacitated supply chains with fixed ordering costs. *Management Science* 48(5): 644-651.

Hammond, J. 1989. Norton auto supply. Harvard Business School Case 9-689-084, 1989, Cambridge, MA.

Kapuscinski, R., and S. Tayur. 1998. A capacitated production-inventory model with periodic demand. *Operations Research* 46(6): 899-911.

Lee, H., K. C. So, and C. S. Tang. 2000. Value of information sharing in a two-level supply chain. *Management Science* 46(5): 626-643.

Lee, H., P. Padmanabhan, and S. Whang. 1996. Information distortion in a supply chain: the bullwhip effect. *Management Science* 43(4): 546-558.

Zhu, W., S. Gavirneni, and R. Kapuscinski. 2007. Improving supply chain performances by imposing simple restrictions while using information flows. Working Paper.

AUTHOR BIOGRAPHIES

LUCY G. CHEN works as an assistant professor in the Business School of National University of Singapore. She received her PHD degree in Aug 2007 from the Johnson Graduate School of Management of Cornell University. Before that, she obtained a BE degree from Tsinghua University, China. Her current research interest lies in the areas of inventory control, (global) supply chain management, the interface on finance and operations management, and scheduling.

SRINAGESH GAVIRNENI teaches operations management at the Johnson Graduate School of Management of Cornell University. Previously he taught at the Kelley School of Business, Indiana University, was the chief algorithm design engineer of SmartOps, a software architect at Maxager Technology, Inc., and a research scientist with Schlumberger. He earned an undergraduate degree in mechanical engineering from IIT–Madras, a masters degree from Iowa State University and a doctorate from Carnegie Mellon University. His research concerns supply chain management, inventory control, production scheduling, simulation, and optimization. He is a frequent contributor to the Winter Simulation Conference and his papers have appeared in many journals including Management Science, Operations Research, Manufacturing & Service Operations Management, IIE Transactions, and Interfaces.