The Value of Information Sharing in a Two-Stage Supply Chain with Production Capacity Constraints

David Simchi-Levi,[1] Yao Zhao[2]

[1] *The Engineering Systems Division and the Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139-4307*

[2] *Department of Management Science and Information Systems, Rutgers University, Newark, New Jersey 07102-1895*

Received 23 May 2001; revised 7 April 2003; accepted 7 April 2003

DOI 10.1002/nav.10094

Abstract: We consider a simple two-stage supply chain with a single retailer facing i.i.d. demand and a single manufacturer with finite production capacity. We analyze the value of information sharing between the retailer and the manufacturer over a finite time horizon. In our model, the manufacturer receives demand information from the retailer even during time periods in which the retailer does not order. To analyze the impact of information sharing, we consider the following three strategies: (1) the retailer does not share demand information with the manufacturer; (2) the retailer does share demand information with the manufacturer and the manufacturer uses the optimal policy to schedule production; (3) the retailer shares demand information with the manufacturer and the manufacturer uses a greedy policy to schedule production. These strategies allow us to study the impact of information sharing on the manufacturer as a function of the production capacity, and the frequency and timing in which demand information is shared. © 2003 Wiley Periodicals, Inc. Naval Research Logistics 50: 888–916, 2003.

Keywords: production-inventory systems; capacity constraint; information sharing; optimal policies; dynamic programming

1. INTRODUCTION

Information technology is an important enabler of efficient supply chain strategies. Indeed, much of the current interest in supply chain management is motivated by the possibilities introduced by the abundance of data and the savings inherent in sophisticated analysis of these data. For example, information technology has changed the way companies interact with suppliers and customers. Strategic partnering, which relies heavily on information sharing, is becoming ubiquitous in many industries.

Correspondence to: D. Simchi-Levi

As observed by Stein and Sweat [9], sharing demand information vertically among supply chain members has achieved huge success in practice. According to Stein and Sweat, by "exchanging information, such as inventory level, forecasting data, and sales trends, these companies are reducing their cycle times, fulfilling orders more quickly, cutting out millions of dollars in excess inventory, and improving customer service."

These developments have motivated the academic community to explore the benefits of information sharing. An excellent review of recent research can be found in Cachon and Fisher [2]. The paper by Aviv and Federgruen [1] is closely related to our work. In their paper, Aviv and Federgruen analyze a single supplier multiple retailer system where retailers face random demand and share inventories and sales data with the supplier. They analyze the effectiveness of a Vendor Managed Inventory (VMI) program where sales and inventory data are used by the supplier to determine the timing and the amount of shipments to the retailers. For this purpose, they compare the performance of the VMI program with that of a traditional, decentralized system, as well as a supply chain in which information is shared continuously, but decisions are made individually, i.e., by the different parties. The objective in the three systems is to minimize the long-run average cost. Aviv and Federgruen report that information sharing reduces system wide cost by 0–5% while VMI reduces cost, relative to information sharing, by 0.4–9.5% and on average by 4.7%. They also show that information sharing could be very beneficial for the supplier.

Thus, the objective of the current paper is not only to characterize the benefit of information sharing, but also to understand how to share information, e.g., how frequently should information be shared and when should it be shared so that the *supplier* can realize the potential benefits. Specifically, our focus, in this paper, is on the so called *Quick Response* strategy (see Simchi-Levi, Kaminsky, and Simchi-Levi [8]), in which demand information is shared continuously but decisions are made by individual parties.

Of course, our paper is not the first one to focus on the potential benefits of information sharing for the supplier. For instance, Gavirneni, Kapuscinski, and Tayur [4] analyzes a simple two-stage supply chain with a single capacitated supplier and a single retailer. In this periodic review model, the retailer makes ordering decisions every period, using an (s, S) inventory policy, and transfers demand information to the supplier every period, independent of whether an order is made. Assuming zero transportation lead time, they show that the benefit, i.e., the supplier's cost savings, due to information sharing, increases with production capacity, and it ranges from 1% to 35%.

In this paper, we investigate a single product, periodic review, two-stage production–inventory system with a single capacitated manufacturer and a single retailer facing i.i.d. demand and using an order-up-to inventory policy. The retailer has a fixed ordering interval. That is, every T time periods, e.g., 4 weeks, the retailer places an order to raise her inventory position to a certain level. The manufacturer receives demand information from the retailer every τ units of time, $\tau \leq T$. For instance, the retailer places an order every 4 weeks but provides demand information every week. This is clearly the case in many retailer–manufacturer partnerships in which orders are placed by the retailer at certain points in time but Point-of-Sale (POS) data are provided every day or every week. In all these cases, POS data is provided to the manufacturer more frequently than the retailers' orders. We refer to the time between successive orders as the *ordering period* and the time between successive information sharing as the *information period.* Of course, in most supply chains, information can be shared almost continuously, e.g., every second, while decisions are made less frequently, e.g., every week. Thus, information periods really refer to the time intervals between successive use of the information provided.

Intuitively, if the retailer shares demand information more frequently than placing orders, the manufacturer can better manage its production and inventory activities. Thus, the manufacturer can reduce her safety stock while maintaining or increasing the service levels. To quantify this intuition, our objective in this paper is to characterize the benefits of information sharing as well as to identify methods that allow the manufacturer to efficiently use this information.

For this purpose, we analyze and compare the following three strategies. In the first strategy, referred to as *no information sharing,* the retailer does not share information with the manufacturer except for order information. In the second strategy, referred to as *information sharing with optimal policy,* the retailer shares demand information with the manufacturer at the end of each information period. We assume that the manufacturer knows the external demand distribution for each information period, and uses an optimal strategy to schedule production so as to minimize her own expected holding and shortage cost. In the third strategy, referred to as *information sharing with greedy policy,* the retailer shares demand information with the manufacturer just as in the previous strategy, but instead of the optimal policy, the manufacturer uses a simple heuristic that is easy to implement, based on demand and shortage in the previous information period, as well as her production capacity.

The rest of this paper is organized as follows: In Section 2, we set up the models for the three strategies, identify the policies used by the manufacturer, discuss their properties, and show the value of information sharing. In Section 3, the optimal timing of information sharing is discussed. In Section 4, we compare the performance of the three strategies using a numerical study. Section 5 concludes the paper and points out the limitations of the paper.

2. MODELS

We consider a periodic review, single product, two-stage system with a single retailer and a single manufacturer. External demand faced by the retailer every information (ordering) period is an i.i.d. random variable. To simplify the analysis, we assume that the retailer controls her inventory position (outstanding orders plus on-hand inventory minus backorders) by an order-up-to policy with a constant order-up-to level; i.e., in every ordering period, the retailer orders to raise her inventory position to the order-up-to level, and this level does not change over time. All unsatisfied demand at the retailer is backlogged; thus the retailer transfers external demand of each ordering period to the manufacturer. The manufacturer has a production capacity limit, i.e., a limit on the amount that the manufacturer can produce per unit of time. The manufacturer runs her production line always at the full capacity limit. Our objective is to compare the performance of the three strategies in a finite time horizon.

The sequence of events in our model is as follows. At the beginning of an ordering period the retailer reviews her inventory and places an order to raise the inventory position to the target inventory level. The manufacturer receives the order from the retailer, fills the order as much as she can from stock, and then makes a production decision. If the manufacturer cannot satisfy all of a retailer's order from stock, then the missing amount is backlogged. The backorder will not be delivered to the retailer until the beginning of the next ordering period. Finally, transportation lead-time between the manufacturer and the retailer is assumed to be zero. Similarly, at the beginning of an information period, the retailer transfers the sales data of previous information period to the manufacturer. Upon receiving this demand information, the manufacturer reduces this demand from her inventory position, although she still holds the stock, and then makes a production decision.

Throughout this paper, we equally divide each ordering period into an integer number of information periods unless otherwise mentioned. Thus, $N = T/\tau$ is an integer, and it represents the number of information periods in one ordering period. We index information periods within one ordering period 1, 2, . . . , N where N is the first information period in the ordering period and 1 is the last. Let C denote the production capacity per information period, τ, while $\bar{C}$ denotes the production capacity per ordering period, T. Hence, $\bar{C} = NC$. Finally, c denotes the production cost per item.

Since we calculate the inventory holding cost for each information period, we let h be the inventory holding cost per unit product per information period. Let $0 < \beta \leq 1$ be the time discount factor for one information period; evidently, one unit of product kept in inventory for n information periods, $n = N, N - 1, \ldots, 1$, will incur a total inventory cost $h_n = h(1 + \beta + \cdots + \beta^{n-1})$. To keep the consistency of notation, let $h_0 = 0$. It is easy to see that the earlier the manufacturer makes a production run in one ordering period, the longer she will carry the inventory, and thus the more holding cost she will have to pay. Penalty cost is charged at the end of each ordering period and thus, let π be the penalty cost per backlogged item per ordering period. We use D to denote the end user demand in one information period, τ. D is assumed to be i.i.d., with $f_D(\cdot)$ ($F_D(\cdot)$) being the pdf (cdf) function and μ being its mean. Finally, ΣD is the total end user demand in one ordering period, T.

2.1. No Information Sharing

Recall that in this strategy, the retailer does not share information with the manufacturer. Since the retailer uses an order-up-to policy with a constant order-up-to level, and all unsatisfied demands are fully backlogged, her order equals to the demand during one ordering periods. Thus, we assume that the manufacturer knows the external demand distribution for one ordering period.

Consider a finite horizon model with M ordering periods and N information periods in each ordering period. Ordering periods are indexed in a reverse order, that is, 0 is the index of the last ordering period in the planning horizon, while $M - 1$ is the index of the first ordering period. The ith information period, $i = 1, 2, \ldots, N$, in ordering period m, $m = 0, 1, \ldots, M - 1$ is referred to as the $mN + i$ information period.

Let $U'_{mN+i}(x)$ be the minimum expected inventory and production costs from period $mN + i$ until the end of the planning horizon, when we start period $mN + i$ with an inventory position x.

It is easy to verify that $W'_{mN+i}(x, y)$, the expected inventory and production cost in the information period $mN + i$, given that the period starts with an inventory position x and produces in that period $y - x$, only depends on i. So we replace W'_{mN+i} with W'_i for the ith information period, and write it as follows:

$$W'_i(x, y) = \begin{cases} c(y - x) + h_{i-1}(y - x), & i = 2, \ldots, N, \\ c(y - x) + E(L(y, \sum D)), & i = 1, \end{cases}$$

where $L(y, \Sigma D) = h_N(y - \Sigma D)^+ + \pi(\Sigma D - y)^+$, and $E(L(y, \Sigma D))$ is the expectation of $L(y, \Sigma D)$ with respect to ΣD. In the very first information period, i.e., in information period NM, a cost of $h_N x^+$ will be charged for the initial inventory position.

Let the salvage cost $U'_0(\cdot) \equiv 0$. If the initial inventory position is zero, then

$$U'_{mN+i}(x) = \begin{cases} \min_{x \leq y \leq x+C}\{W'_i(x, y) + \beta U'_{mN+i-1}(y)\}, & i = 2, \ldots, N \quad \text{and} \quad \forall\, m, \\ \min_{x \leq y \leq x+C}\{W'_i(x, y) + \beta E(U'_{mN+i-1}(y - \sum D))\}, & i = 1 \quad \text{and} \quad \forall\, m. \end{cases}$$

To find the optimal policy, for $m = 0, \ldots, M - 1$, we rewrite the dynamic program as follows:

$$U'_{mN+i}(x) = -(c + h_{i-1})x + V'_{mN+i}(x), \qquad \forall\, i,$$

$$V'_{mN+i}(x) = \min_{x \le y \le x+C}\{J'_{mN+i}(y)\}, \qquad \forall\, i,$$

$$J'_{mN+i}(y) = \begin{cases} cy + h_{i-1}y + \beta U'_{mN+i-1}(y), & i = 2, \ldots, N, \\ cy + E(L(y, \sum D)) + \beta E(U'_{mN+i-1}(y - \sum D)), & i = 1. \end{cases}$$

We now discuss properties of the above dynamic program. A straightforward analysis of the finite planning horizon (see Federgruen and Zipkin [3]) shows the following two results:

LEMMA 1: The set $A \equiv \{(x, y)|x \le y \le x + C\}$ is convex. For all $m = 0, \ldots, M - 1$ and $i = 1, \ldots, N$ we have:

(a) $E(L(y, \Sigma D))$, $J'_{mN+i}(y)$, $V'_{mN+i}(x)$, and $U'_{mN+i}(x)$ are convex,
(b) $U'_{mN+i}(x) \to +\infty$, when $|x| \to +\infty$, and
(c) if $\beta^{N-1}\pi > c + h_{N-1}$, then $J'_{mN+i}(y) \to +\infty$ when $|y| \to +\infty$.

See Appendix A for a proof.

LEMMA 2: Let y^*_{mN+i} be the smallest value minimizing J'_{mN+i}, and let x be the inventory position at the beginning of period $mN + i$. Then, y^*_{mN+i} is finite, and the optimal production-inventory policy is to produce

$$\begin{cases} 0, & x \ge y^*_{mN+i}, \\ y^*_{mN+i} - x, & 0 \le y^*_{mN+i} - x \le C, \\ C, & \text{otherwise.} \end{cases}$$

A third, quite intuitive property, is that, given two policies that produce the same amount in a given ordering period, a cost-effective policy will postpone production as much as possible during the ordering period. Of course, this property does not need any proof. We use dynamic programming method to solve for y^*_m in single and multiple ordering period cases.

2.2. Information Sharing with Optimal Policy

In this strategy, the retailer provides the manufacturer with demand information every information period, and the data are used by the manufacturer to optimize her production and inventory costs. We consider the following two cases:

2.2.1. One Ordering Period

We start by considering a single ordering period with N information periods. We follow the convention that N is the first information period and 1 is the last information period. Let I_n be the manufacturer's *on-hand inventory level* at the beginning of the nth information period. D_n represents the demand during the nth information period. We use $x_n \equiv I_n - \Sigma_{i=n+1}^{N} D_i$. Thus,

x_n is the *inventory position* at the beginning of the nth information period. Let y_n be the inventory position at the end of nth information period after production in this period but not taking D_n into account. That is, y_n is equal to x_n plus the amount produced in the nth time period.

Let $U_n(x_n)$ be the minimum total inventory and production costs from the beginning of nth information period until the end of the planning horizon, given an initial inventory position x_n. To simplify the notation, we drop the index n from x_n, y_n, and D_n; this will cause no confusion. Clearly,

$$U_1(x) = \min_{x \le y \le x+C}\{c(y-x) + E(L(y, D))\},$$

$$U_n(x) = \min_{x \le y \le x+C}\{c(y-x) + h_{n-1}(y-x) + \beta E(U_{n-1}(y-D))\}, \qquad n = 2, \dots, N-1,$$

$$U_N(x) = \min_{x \le y \le x+C}\{c(y-x) + h_{N-1}(y-x) + \beta E(U_{N-1}(y-D))\} + h_N x^+. \tag{1}$$

$L(y, D) = h_N(y - D)^+ + \pi(D - y)^+$ and $E(\cdot)$ is the expectation with respect to D, the demand in one information period. Observe that the holding cost for $y - x$ items produced in the nth information period is $h_{n-1}(y - x)$, since these items are kept in inventory from the end of period n until the end of period 1.

Rearranging the equations above, we obtain

$$U_1(x) = -cx + V_1(x),$$

$$V_1(x) = \min_{x \le y \le x+C}\{J_1(y)\},$$

$$J_1(y) = cy + E(L(y, D)),$$

$$U_n(x) = -(c + h_{n-1})x + V_n(x),$$

$$V_n(x) = \min_{x \le y \le x+C}\{J_n(y)\},$$

$$J_n(y) = cy + h_{n-1}y + \beta E(U_{n-1}(y - D)),$$

$$n = 2, \dots, N-1,$$

$$U_N(x) = -(c + h_{N-1})x + h_N x^+ + V_N(x),$$

$$V_N(x) = \min_{x \le y \le x+C}\{J_N(y)\},$$

$$J_N(y) = cy + h_{N-1}y + \beta E(U_{N-1}(y - D)). \tag{2}$$

2.2.2. *Multiple Ordering Periods*

Using the same notation as in the no information sharing model, it is easy to verify that W_i, the expected inventory and production cost in the information period $mN + i$, given that the period starts with an inventory position x and produces in that period $y - x$, can be written as follows:

$$W_i(x, y) = \phi_i(x) + \varphi_i(y), \tag{3}$$

where

$$\phi_i(x) = \begin{cases} -cx, & i = 1, \\ -(c + h_{i-1})x, & \text{otherwise.} \end{cases}$$

$$\varphi_i(y) = \begin{cases} cy + E(L(y, D)), & i = 1, \\ (c + h_{i-1})y, & \text{otherwise,} \end{cases}$$

and

$$L(y, D) = h_N(y - D)^+ + \pi(D - y)^+.$$

Thus, the following recursive relation must hold:

$$U_{mN+i}(x) = \min_{x \le y \le x+C}\{W_i(x, y) + \beta E(U_{mN+i-1}(y - D))\},$$

which can be written as

$$U_{mN+i}(x) = \phi_i(x) + V_{mN+i}(x),$$

$$V_{mN+i}(x) = \min_{x \le y \le x+C}\{J_{mN+i}(y)\},$$

$$J_{mN+i}(y) = \varphi_i(y) + \beta E(U_{mN+i-1}(y - D)). \tag{4}$$

Of course, in the very first information period of the whole planning horizon, we have to add $h_N x^+$ to $U_{MN}(x)$ to account for the holding cost for initial inventory. This is identical to what we did in the no information sharing model and the information sharing model with one ordering period.

Similar properties to Lemmas 1 and 2 can be shown for this model. Specifically,

LEMMA 3: The set $A \equiv \{(x, y)|x \le y \le x + C\}$ is convex. For all $m = 0, \ldots, M - 1$ and $i = 1, \ldots, N$ we have:

(a) $E(L(y, D))$, $J_{mN+i}(y)$, $V_{mN+i}(x)$, and $U_{mN+i}(x)$ are convex,
(b) $U_{mN+i}(x) \to +\infty$, when $|x| \to +\infty$, and
(c) if $\beta^{N-1}\pi > c + h_{N-1}$, then $J_{mN+i}(y) \to +\infty$ when $|y| \to +\infty$.

See Appendix B for a proof.

LEMMA 4: Let y^*_{mN+i} be the smallest value minimizing J_{mN+i}, and let x be the inventory position at the beginning of period $mN + i$. Then, y^*_{mN+i} is finite and the optimal production-inventory policy is to produce

$$\begin{cases} 0, & x \geq y^*_{mN+i}, \\ y^*_{mN+i} - x, & 0 \leq y^*_{mN+i} - x \leq C, \\ C, & \text{otherwise.} \end{cases}$$

The question is whether one can identify the relationship between the optimal order-up-to-levels of two consecutive information periods. Intuitively, delaying production as long as possible within one ordering period should reduce inventory holding cost. The risk, of course, is that delaying too much may lead to a shortage, due to the limited production capacity. Thus, the next property characterizes sufficient conditions under which postponing production is profitable.

PROPOSITION 5: If $Pr(D > C) = 0$, then $y^*_{mN+i} \leq y^*_{mN+i-1}$, for $i = 2, \ldots, N$, $m = 0, 1, \ldots, M - 1$.

PROOF: We prove the result for the last ordering period, i.e., $m = 0$. For $n = 2, \ldots, N$, rewrite Eq. (2) as follows:

$$J_n(y_n) = (1 - \beta)cy_n + (h_{n-1} - \beta h_{n-2})y_n + \beta(c + h_{n-2})E(D) + \beta Q_n(y_n),$$

$$Q_n(y_n) = E(V_{n-1}(y_n - D))$$

$$= E\{\min_{y_n - D \leq y_{n-1} \leq y_n - D + C}[J_{n-1}(y_{n-1})]\}.$$

Let $\bar{y}_n$ be the smallest value minimizing $Q_n(y_n)$. Observe that y^*_{n-1}, the minimizer of $J_{n-1}(y)$, satisfies

$$Q_n(y^*_{n-1}) = J_{n-1}(y^*_{n-1}).$$

This is true since $Pr(D > C) = 0$ and $D \geq 0$, which implies that y^*_{n-1} is feasible, i.e.,

$$y^*_{n-1} - D \leq y^*_{n-1} \leq y^*_{n-1} - D + C,$$

for all realization of D. Hence, $\bar{y}_n \leq y^*_{n-1}$. Furthermore, we notice that the difference between $J_n(y_n)$ and $Q_n(y_n)$ is a linearly increasing function, so the first-order right-hand derivative of J_n is positive at y^*_{n-1} (the first-order right-hand derivative exists for J_n because J_n is convex). Finally, since J_n is convex, the result follows. The proof for all the other ordering periods is identical. □

In practice, of course, the assumption that $Pr(D > C) = 0$ may not always hold, and thus the question is whether one can identify other situations where we can characterize the relationship between y^*_n and y^*_{n-1}.

Observe that if $Pr(D > C) > 0$, then $\bar{y}_n \geq y^*_{n-1}$, since $Q_n(y_n) \geq Q_n(y^*_{n-1})$ for $y_n < y^*_{n-1}$. Thus, a result similar to Proposition 5 can not be proven in the same way.

Because $J_n(y_n)$, $n = 1, \ldots, N$ are convex, they are continuous and right-hand differentiable. Hence, define $\Delta = d/dy$ to be the right-hand derivative, the following equations hold:

$$\Delta J_n(y_n) = (1 - \beta)c + (h_{n-1} - \beta h_{n-2}) + \Delta\beta Q_n(y_n)$$

$$= (1 - \beta)c + (h_{n-1} - \beta h_{n-2}) + \beta \int_0^{(y_n - y^*_{n-1})^+} \Delta J_{n-1}(y_n - D) f_D(D)\, dD$$

$$+\beta \int_{(y_n - y^*_{n-1})^+}^{\infty} \Delta J_{n-1}(y_n - D) f_D(D + C)\, dD.$$

Clearly, if $\Delta J_n(y^*_{n-1}) \geq 0$, then from the convexity and the limiting behavior of J_n, we have $y^*_n \leq y^*_{n-1}$. Thus, plug in y^*_{n-1}:

$$\Delta J_n(y^*_{n-1}) = (1 - \beta)c + (h_{n-1} - \beta h_{n-2}) + \beta \int_0^{\infty} \Delta J_{n-1}(y^*_{n-1} - D) f_D(D + C)\, dD.$$

Since $\Delta J_{n-1}(y^*_{n-1} - D) \leq 0$ for $D \geq 0$, it is not clear whether $\Delta J_n(y^*_{n-1}) \geq 0$. We use numerical methods to evaluate ΔJ_{n-1} in our computational study.

2.3. The Value of Information Sharing

In this subsection we quantify the benefits from information sharing in a model with M ordering periods and N information periods in each ordering period. Our focus is on the extreme case in which production capacity is infinite so that the manufacturer only needs to produce in the last information period. Notice that the sequence of events in our model excludes a make-to-order policy when production capacity is infinite. That is, in our model, the manufacturer will satisfy the order only from her on-hand stock. If the manufacturer does not have enough stock on hand, she will pay a penalty cost for backlogging the missing amount. Evidently, this is the limiting case of the information sharing model (Section 2.2) as the production capacity approaches infinity.

First, let us consider the no information sharing strategy. The cost function for the last ordering period is

$$B'_1(x) = c(y - x) + L(y, \sum D_1) = -cx + g(y, \sum D_1),$$

where x is the initial inventory position at the beginning of the ordering period, y is the target inventory position, ΣD_k is the total demand in the kth ordering period and

$$g(y, \sum D_1) = cy + L(y, \sum D_1).$$

Let $\alpha = \beta^N$ be the time discount factor for one ordering period. Since salvage cost is equal to zero, the total cost in M ordering periods is

$$B'_M(x_M) = h_N x_M^+ + \sum_{m=0}^{M-1} \alpha^m[-cx_{M-m} + g(y_{M-m}, \sum D_{M-m})],$$

given that the initial inventory position of the planning horizon is x_M.

Since $x_m = y_{m+1} - \sum D_{m+1}$ for $m = M - 1, \ldots, 1$, a straightforward calculation shows that

$$E(B'_M(x_M)) = h_N x_M^+ - cx_M$$
$$+ E\left[\sum_{m=0}^{M-2} \alpha^m(g(y_{M-m}, \sum D_{M-m}) - c\alpha y_{M-m}) + \alpha^{M-1}g(y_1, \sum D_1) + \sum_{m=0}^{M-1} \alpha^m c \sum D_{M-m}\right],$$

where $E(\cdot)$ is the expectation with respect to demand $\sum D_m$, $m = M, M - 1, \ldots, 1$.

In what follows, we omit $h_N x_M^+$ by setting $x_M = 0$ because it is the same for both the information sharing model and the no information sharing model. Since our focus is on the impact of information sharing on the inventory costs, we ignore production cost in our model. Hence,

$$E(B'_M(x_M)) = E\left[\sum_{m=0}^{M-1} \alpha^m L(y_{M-m}, \sum D_{M-m})\right],$$

where $y_m \geq x_m$ for $m = M, M - 1, \ldots, 1$.

To simplify the model, we assume that demand has independent and identical increments. Define D_t to be the demand in any time period of length t; thus $D_T = \sum D$ and $D_\tau = D$. To simplify the notation, we let $G(y, t) = E(L(y, D_t))$. Following Heyman and Sobel [6], it can be shown that if $Prob\{D_T < 0\} = 0$, myopic policy is optimal for the manufacturer. Further, let y_T^* be the optimal order-up-to level for the myopic policy; if the initial inventory position $x_M \leq y_T^*$, then $U'_{MN}(x_M)$, the minimum expected inventory cost from the information period MN to the end of the planning horizon satisfies

$$U'_{MN}(x_M) = \frac{1 - \alpha^M}{1 - \alpha} G(y_T^*, T).$$

In order to obtain analytic result, we further assume that demand D_t can be approximated by $Normal(t\mu, t\sigma^2)$. Notice that in this case $Prob\{D_t < 0\} > 0$. One way to constrain the probability of this happening is to choose μ and σ so that $Prob\{D_t < 0\} \leq \epsilon$, where $\epsilon > 0$ is a very small number.

Let $\Phi(\cdot)$ be the standard normal cumulative distribution function, and $\phi(\cdot)$ be the standard normal density function. Hence,

$$G(y, t) = h_N \int_{-\infty}^{y} (y - \xi) f_{D_t}(\xi)\, d\xi + \pi \int_{y}^{\infty} (\xi - y) f_{D_t}(\xi)\, d\xi$$

$$= \pi(t\mu - y) + (h_N + \pi) \int_{-\infty}^{y} (y - \xi) f_{D_t}(\xi)\, d\xi.$$

We denote $\gamma = \pi/(\pi + h_N)$, and $z_\gamma = \Phi^{-1}(\gamma)$. From the analysis of the celebrated news vendor problem, we know that $G(y, t)$ reaches its minimum at $y_t^* = t\mu + z_\gamma\sqrt{t}\,\sigma$. Let $\eta = (\xi - t\mu)/\sqrt{t}\,\sigma$; hence,

$$G(y_t^*, t) = \pi t\mu - (h_N + \pi)\int_{-\infty}^{z_\gamma} (t\mu + \sqrt{t}\,\sigma\eta)\phi(\eta)\,d\eta$$

$$= (h_N + \pi)\sqrt{t}\,\sigma\kappa,$$

where $\kappa = -\int_{-\infty}^{z_\gamma} \eta\phi(\eta)\,d\eta$.

Next, consider the information sharing strategy. The cost function for one ordering period is

$$c(y - x + D_{T-\tau}) + L(y, D_\tau),$$

where x is defined in the same way as in the no information sharing strategy, y is the target inventory position of the last information period by taking $D_{T-\tau}$ into account, $D_{T-\tau}$ is the realized demand in information periods $N, N - 1, \ldots, 2$, and D_τ is the demand in the last information period. That is, $D_{T-\tau} + D_\tau$ is the demand realized in this ordering period. For simplicity, let $D' = D_{T-\tau}$ and $D = D_\tau$.

Assuming zero production cost and following the same procedure as in the no information sharing model, we have

$$E(B_M(x_M)) = E\left[\sum_{m=0}^{M-1} \alpha^m L(y_{M-m}, D_{M-m})\right],$$

with $y_m \geq x_m - D'_m$ for $m = M, M - 1, \ldots, 1$.

Thus, if the initial inventory position $x_M \leq y_\tau^*$, $U_{MN}(x_M) = [(1 - \alpha^M)/(1 - \alpha)]G(y_\tau^*, \tau)$. These results lead to the following observations for the model with infinite production capacity:

- The information sharing strategy has the same fill rate as the no information sharing strategy.
- The expected cost in the information sharing strategy is proportional to $\sqrt{\tau}$ while the expected cost under no information sharing is proportional to $\sqrt{T} = \sqrt{N\tau}$, where N is the number of information periods in one ordering period. Thus, the percentage cost saving due to information sharing (defined to be the ratio of cost saving due to information sharing to the cost of no information sharing) is proportional to $1 - \sqrt{(1/N)}$. For example, sharing information 4 times in one ordering period allows the manufacturer to reduce total inventory cost by 50% relative to no information sharing.

2.4. Analysis of Nondimensional Parameters

Our objective in this section is to identify the parameters which may have an impact on the percentage cost reduction due to information sharing. For simplicity, we focus on a single ordering period, but a similar method can be applied to the problem with any number of ordering period.

Divide both sides of Eq. (1) by $h_N N\mu$, we obtain

$$\frac{U_1(x)}{h_N N\mu} = \min_{(x/N\mu)\le(y/N\mu)\le(x/N\mu)+(1/N)(C/\mu)} \left\{\frac{c}{h_N}\frac{(y-x)}{N\mu} + \frac{1}{h_N N\mu}\int L(y,\xi) f_D(\xi)\, d\xi\right\}.$$

Let $\eta = \xi/N\mu$, $D' = D/N\mu$; it is easy to see that

$$f_D(\xi) = \frac{d}{d\xi}F_D(\xi) = \frac{d}{d\xi}Pr\left(\frac{D}{N\mu} \le \frac{\xi}{N\mu}\right) = \frac{d}{d\xi}F_{D'}\left(\frac{\xi}{N\mu}\right) = \frac{1}{N\mu}f_{D'}(\eta).$$

Hence,

$$\frac{1}{h_N N\mu}\int L(y,\xi) f_D(\xi)\, d\xi = \int \left(\left(\frac{y}{N\mu} - \eta\right)^+ + \frac{\pi}{h_N}\left(\eta - \frac{y}{N\mu}\right)^+\right) f_{D'}(\eta)\, d\eta.$$

Let $x' = x/N\mu$, $y' = y/N\mu$, $c' = c/h_N$, $\rho = \mu/C$, $\pi' = \pi/h_N$, $U_1' = U_1/h_N N\mu$, and $L'(y', \eta) = (y' - \eta)^+ + \pi'(\eta - y')^+$. We omit $'$ from the notation without causing any confusion, and we can rewrite the nondimensionalized function U_1 as follows:

$$U_1(x) = \min_{x\le y\le x+(1/N)(1/\rho)} \{c(y-x) + \int L(y,\eta) f_D(\eta)\, d\eta\}.$$

A similar technique can be applied to $U_n(x)$ for $n = 2, \ldots, N$. Hence,

$$U_n(x) = \min_{x\le y\le x+(1/N)(1/\rho)}\left\{c(y-x) + \frac{h_{n-1}}{h_N}(y-x) + \beta E(U_{n-1}(y-D))\right\},$$

$$n = N-1, \ldots, 2,$$

$$U_N(x) = \min_{x\le y\le x+(1/N)(1/\rho)}\left\{c(y-x) + \frac{h_{N-1}}{h_N}(y-x) + \beta E(U_{N-1}(y-D))\right\} + x^+.$$

Thus, the percentage cost reduction associated with information sharing relative to no information sharing depends only on the following nondimensional parameters: ρ, N, π, c, β, and $f_D(\eta)$, where ρ is the capacity utilization μ/C, N is the frequency of information sharing, π and c are the nondimensionalized penalty and production costs, and $f_D(\eta)$ is the probability density function of the nondimensionalized demand. In our computational study, we will focus on the impact of these parameters on the benefit from information sharing.

2.5. Information Sharing and the Greedy Policy

In this strategy, we apply a simple heuristic that specifies production quantities so as to match supply and demand. This heuristic is motivated by our experience with a number of manufacturing companies who use the shared information in a greedy fashion. That is, these companies use a "lot-by-lot" strategy; i.e., they continuously update production quantities to match realized demand. Thus, our objective is to quantify the impact of the greedy heuristic on the manufacturer's cost and service level.

Specifically, in this heuristic the manufacturer produces in every information period, $n = N - 1, N - 2, \ldots, 2$, an amount equal to

$$\min\{C, D_{n+1} - x_n^-\},$$

where $x_n^- = \min\{0, x_n\}$ is the shortage level at the beginning of the nth information period. In the first information period, i.e., $n = N$, the manufacturer produces

$$\min\{C, -x_N^-\}.$$

Finally, in the last information period, i.e., $n = 1$, the inventory is raised to a certain level determined by production capacity, inventory at the beginning of the information period, production and inventory holding costs, and the demand distribution. This can be done by solving a newsboy problem with capacity constraint; the reader is referred to Hadley and Whitin [5] or Lee and Nahmias [7] for a review of the newsboy problem.

3. TIMING OF INFORMATION SHARING

In this section, we fix the frequency of information sharing, i.e., N is fixed, but allow the times at which information is shared to change. Thus, an important question is when to share information. To simplify the analysis, we focus on the single ordering period model and assume that the retailer can share information with the manufacturer only once during the ordering period. Intuitively, the higher the production capacity per unit time, the later the time information may be shared. Of course, the later the timing of information sharing, the more accurate the information on demand during the ordering period but the smaller the remaining production capacity, i.e., the product of the per unit time production capacity and the remaining time until the end of the ordering period. For instance, if production capacity per unit of time is very high, information should be transferred and used almost at the end of the ordering period. On the other hand, if production capacity is tight, then it is less clear when information should be shared. Thus, our objective is to find (1) the optimal time to share information, and (2) the parameters which may affect the best timing for sharing information.

In order to find the optimal timing, we develop a continuous time model. For this purpose, all notations associated with an information period will be changed to per unit of time, while the others remain the same. Hence, h is the inventory holding cost per unit product per unit of time; C is the production capacity per unit of time; and D is the customer demand per unit of time with mean μ. Finally, to simplify the analysis, we set $\beta = 1$. Similar to the discrete time model, production is delayed as long as possible until the time that the target inventory position can just be reached by producing at the rate C.

Consider an ordering interval $[0, T]$ and a given $t < T$, let $T - t$ be the time when information is shared. $t = 0$ ($t = T$) implies that information is shared at the end (beginning) of the ordering period. We assume that customer demand D_τ in any time interval of length τ is $Poisson(\tau\mu)$. This implies that customer demand at any time interval $[t, t + \tau] \subset [0, T]$ of length τ depends only on τ but not on t, and demand in different (not overlapping) time intervals is independent. The dynamic program is formulated as below. Given that information is shared after $T - t$ units of time, let $U_1(x, t)$ ($U_2(x, t)$) be the minimum expected inventory and production costs from the time information is shared (the time the ordering period starts, respectively) to the end of the horizon given an initial inventory position x:

$$U_2(x, t) = \min_{x \le y \le x+(T-t)C}\{c(y - x) + H_2(x, y, t) + E(U_1(y - D_{T-t}, t))\},$$

$$U_1(x, t) = \min_{x \le y \le x+tC}\{c(y - x) + H_1(x, y) + E(L(y, D_t))\}, \tag{6}$$

where

$$H_2(x, y, t) = T \times h \times x^+ + h \times t \times (y - x) + \frac{h}{2}\frac{(y - x)^2}{C},$$

$$H_1(x, y) = \frac{h}{2}\frac{(y - x)^2}{C},$$

$$L(y, D) = T \times h \times (y - D)^+ + \pi \times (D - y)^+.$$

To show that we correctly calculate the total inventory holding cost, let us assume that x_1 is the initial inventory position at the beginning of the ordering period, y_1 is the order-up-to level at the end of the first information period, $x_2 = y_1 - D_{T-t}$ is the inventory position at the beginning of the second information period, and y_2 is the order-up-to level at the end of the second information period. With these notations, the total inventory holding cost during the ordering period equals to

$$h \times T \times x_1^+ + h \times t \times (y_1 - x_1) + \frac{h}{2}\frac{(y_1 - x_1)^2}{C} + \frac{h}{2}\frac{(y_2 - x_2)^2}{C}.$$

Notice that the term

$$h \times T \times x_1^+ + h \times t \times (y_1 - x_1) + \frac{h}{2}\frac{(y_1 - x_1)^2}{C}$$

is a function of x_1, y_1, and t only and hence it is captured by $H_2(x, y, t)$. The term

$$\frac{h}{2}\frac{(y_2 - x_2)^2}{C}$$

is a function of x_2 and y_2 only and hence is captured by $H_1(x, y)$.

If t is fixed, it is easy to show that $V_1(x, y, t) = c(y - x) + H_1(x, y) + E(L(y, D_t))$ and $V_2(x, y, t) = c(y - x) + H_2(x, y, t) + E(U_1(y - D_{T-t}, t))$ are jointly convex in both x

and y [since the Hessians of $H_1(x, y)$, $H_2(x, y, t)$ are positive semidefinite]. This observation implies:

PROPOSITION 6: $U_1(x, t)$ and $U_2(x, t)$ are convex in x.

PROOF: We start by proving that $U_1(x, t)$ is convex in x. Suppose we have x_1, x_2, $x_1 \neq x_2$, and y_1^*, y_2^*, where $U_1(x_1, t) = V_1(x_1, y_1^*, t)$ with $x_1 \leq y_1^* \leq x_1 + tC$ and $U_1(x_2, t) = V_1(x_2, y_2^*, t)$ with $x_2 \leq y_2^* \leq x_2 + tC$. Let $\bar{x} = \lambda x_1 + (1 - \lambda)x_2$ and $\bar{y} = \lambda y_1^* + (1 - \lambda)y_2^*$, obviously, for any $\lambda \in (0, 1)$, $\bar{x} \leq \bar{y} \leq \bar{x} + tC$. Hence,

$$
\begin{aligned}
U_1(\bar{x}, t) &= \min\{V_1(\bar{x}, y, t) | \bar{x} \leq y \leq \bar{x} + tC\} \\
&\leq V_1(\bar{x}, \bar{y}, t) \\
&\leq \lambda V_1(x_1, y_1^*, t) + (1 - \lambda)V_1(x_2, y_2^*, t) \\
&= \lambda U_1(x_1, t) + (1 - \lambda)U_1(x_2, t).
\end{aligned}
$$

To prove that $U_2(x, t)$ is convex in x, observe that since $U_1(x, t)$ is convex in x, thus

$$V_2(x, y, t) = c(y - x) + H_2(x, y, t) + E(U_1(y - D_{T-t}, t))$$

is jointly convex in x and y. Applying the same proof as before, we can show that $U_2(x, t)$ is convex in x. □

Unfortunately, it is not clear whether or not $E(U_1(y - D_{T-t}, t))$ and $U_2(x, t)$ are convex in t. Thus, for any given t, we compute the optimal y by utilizing Proposition 6. To find the optimal timing of information sharing, we discretize the ordering period and search on all possible values of t.

To characterize the optimal timings of information sharing when the retailer can transfer demand information to the manufacturer more than once, we extend our analysis to the case in which the retailer can share information with the manufacturer twice during one ordering period. Consider an ordering period $[0, T]$; let $0 \leq t_1 \leq t_2 \leq T$ be the times when information is shared, $U_1(x, t_2)$ ($U_2(x, t_1, t_2)$, $U_3(x, t_1, t_2)$) be the minimum expected production and inventory cost from the time of the second information sharing (the time of the first information sharing, the time the ordering period starts, respectively) to the end of the ordering period given an initial inventory position x. Then the following dynamic program holds:

$$U_3(x, t_1, t_2) = \min_{x \leq y \leq x + t_1 C}\{c(y - x) + K_3(x, y, t_1) + E(U_2(y - D_{t_1}, t_1, t_2))\},$$

$$U_2(x, t_1, t_2) = \min_{x \leq y \leq x + (t_2 - t_1)C}\{c(y - x) + K_2(x, y, t_2) + E(U_1(y - D_{t_2 - t_1}, t_2))\},$$

$$U_1(x, t_2) = \min_{x \leq y \leq x + (T - t_2)C}\{c(y - x) + K_1(x, y) + E(L(y, D_{T - t_2}))\},$$

where

$$K_3(x, y, t_1) = T \times h \times x^+ + \frac{h}{2}\frac{(y-x)^2}{C} + (T - t_1) \times h \times (y - x),$$

$$K_2(x, y, t_2) = \frac{h}{2}\frac{(y-x)^2}{C} + (T - t_2) \times h \times (y - x),$$

and

$$K_1(x, y) = \frac{h}{2}\frac{(y-x)^2}{C}.$$

Following the same procedure as in Proposition 6, we can show that $U_3(x, t_1, t_2)$, $U_2(x, t_1, t_2)$, and $U_1(x, t_2)$ are convex in x. Since it is not clear whether $U_3(x, t_1, t_2)$ is convex in t_1 and t_2, we compute the optimal timing t_1 and t_2 by discretizing the ordering period and searching on all possible values of t_1 and t_2.

Finally, we identify the nondimensional parameters which may have impacts on the optimal timing(s). We start with the case in which information is shared only once in one ordering period. We divide Eq. (6) by $hT^2\mu$; let $x' = x/T\mu$, $y' = y/T\mu$, $t' = t/T$, $c' = c/Th$, $\rho = \mu/C$, $\pi' = \pi/Th$, $U_1' = U_1/hT^2\mu$, $U_2' = U_2/hT^2\mu$, $H_1' = H_1/hT^2\mu$, $H_2' = H_2/hT^2\mu$, $D' = D/T\mu$, and $L'(y', \eta) = (y' - \eta)^+ + \pi'(\eta - y')^+$. We omit $'$ from the notation without creating any confusion, and rewrite the nondimensionalized functions $U_1(x, t)$ and $U_2(x, t)$ as follows:

$$U_2(x, t) = \min_{x \le y \le x + [(1-t)/\rho]}\{c(y - x) + H_2(x, y, t) + E(U_1(y - D_{1-t}, t))\},$$

$$U_1(x, t) = \min_{x \le y \le x + (t/\rho)}\{c(y - x) + H_1(x, y) + E(L(y, D_t))\},$$

where $H_2(x, y, t) = x^+ + t(y - x) + (\rho/2)(y - x)^2$, $H_1(x, y) = (\rho/2)(y - x)^2$, and $L(y, D) = (y - D)^+ + \pi(D - y)^+$. This analysis shows that the nondimensional optimal timing is a function of ρ, π, c, and $f_D(\cdot)$ only, of which we will study the effects of ρ and π in the following section.

In the case when information can be shared twice in one ordering period, a similar nondimensional analysis shows that the nondimensional optimal timings t_1/T, t_2/T only depend on ρ, π, c, and $f_D(\cdot)$.

4. COMPUTATIONAL RESULTS

In this section, we use computational analysis to develop insights on the benefits of information sharing. Our goal is twofold: (1) Determine situations in which information sharing provides significant cost savings compared to supply chains with no information sharing; (2) identify the benefits of using information optimally compared to using information greedily. Our focus is on the manufacturer's cost and service level.

According to Section 2.4, we examined cases with variations on the following nondimensional parameters: production capacity over mean demand, the number of information periods in one ordering period, the time when information is shared, demand distribution, and finally the ratio between penalty cost and inventory holding cost.

We set production cost equal to zero, and focus on holding and penalty costs. Let inventory holding cost per ordering period to be equal to a constant 0.4 $ per unit product for all cases. Thus, the inventory holding cost per information period is $0.4/N$, where N is the number of information periods within one ordering period. Penalty cost varies from 1.9 to 7.9 $ per unit product per ordering period and takes the following values 1.9, 3.4, 4.9, 6.4, 7.9. In all cases of our computational analysis, the time discount factor β is assumed to be 1.

Let the initial inventory position at the beginning of the first ordering period, x, be equal to zero. To simplify the calculation, we use discrete probability distributions for customer demand in one information period. In our study, we consider discrete distributions such as Poisson, Uniform, and Binomial. In addition, we also analyze the following discrete distributions: the first, referred to as Disc1, demand takes values from the set (0, 1, 3, 6) with probability (0.1, 0.3, 0.5, 0.1), respectively. The second, referred to as Disc2, demand takes the same values with probability (0.05, 0.2, 0.7, 0.05), respectively.

The dynamic program algorithms such as value iteration allows us to find the cost associated with the first two strategies. For the third strategy, the newsboy model allows us to find the optimal order-up-to level in the last information period of every ordering period, while the cost associated with the strategy is estimated through simulation. Finally simulation results provide us with service levels for all three strategies. Following convention, we measure service level by the type one fill rate, which is defined to be the fraction of ordering periods in which no backorders occur.

In the simulation models, each system is simulated 40,000 times. The fill rate is calculated as follows: Let X_i be a random variable taking the value one if demand (at the end of the ordering period) is satisfied with no shortage in the ith run, and zero otherwise. Our estimation of the type one fill rate is the sample mean $\bar{X} = \sum_i X_i/n$, where n is 40,000. Since our estimation of the standard deviation of X_i is equal to $\sqrt{\bar{X}(1 - \bar{X})}$, which is less than 0.5, thus, the length of a 95% confidence interval of the fill rate is no more than 0.0098.

The following discussions are based on our computational results for models with one ordering period and four information periods unless otherwise mentioned. For multiordering period planning horizon, similar results are obtained.

4.1. The Effect of Information Sharing on the Optimal Policy

In this subsection we analyze the impact of production capacity, penalty cost and demand variability on the optimal policy when information is shared.

Table 1 presents the effect of production capacity for three different distributions of demand in one information period. For each demand distribution, we increase the ratio of production capacity to mean demand (the column capacity/mean demand) and calculate the order-up-to-levels in all information periods. Thus, the last column represents the order-up-to-level for each of the four information periods.

Observe that

- Proposition 2.5 holds for almost all cases except the one in which production capacity is very tight, e.g., capacity/mean demand = 1.2.
- As capacity increases, the difference between order-up-to levels in different information periods increases. The intuition is clear: as capacity increases, the optimal policy delays production as much as possible.
- The order-up-to-levels in the first few information periods may be negative, which implies that the inventory position can be negative.

Table 1. The impact of production capacity.

Demand Distribution	Capacity/Mean Demand	Penalty/Holding Costs	Order-up-to-Levels
Poisson(5)	1.2	8.5	(8, 9, 9, 8)
Poisson(5)	1.6	8.5	(2, 5, 7, 8)
Poisson(5)	2	8.5	(−4, 1, 5, 8)
Uniform(0, 1, . . . , 9)	1.22	8.5	(10, 10, 10, 8)
Uniform(0, 1, . . . , 9)	1.67	8.5	(4, 6, 8, 8)
Uniform(0, 1, . . . , 9)	2.11	8.5	(−2, 2, 6, 8)
Binomial(0.5, 10)	1.2	8.5	(6, 7, 7, 7)
Binomial(0.5, 10)	1.6	8.5	(0, 3, 5, 7)
Binomial(0.5, 10)	2	8.5	(−6, −1, 3, 7)

Table 2 demonstrates the impact of penalty cost. In this table, we increase the ratio of penalty to holding costs from 4.75 to 19.75 for each demand distribution. The table demonstrates the following insights:

- As penalty cost increases, the order-up-to-levels increase.
- As penalty cost increases, the differences between consecutive order-up-to levels decreases.

Table 3 presents the impact of demand variability. In this case the capacity over average demand was kept constant, at a level of 1.67 for all cases, and penalty over holding cost was 7.9 for all cases. It is easy to see that the coefficient of variation of demand distribution has a similar impact as the penalty cost. That is:

- As the coefficient of variation increases, the order-up-to-levels increase.
- As the coefficient of variation increases, the differences between consecutive order-up-to levels decrease.

In Table 4 we consider two ordering periods with four information periods in each one. We observe that the difference between the order-up-to-levels of the same information periods in different ordering periods is small relative to the average total demand in one ordering period. For example, in the case of the Binomial demand distribution, the maximum difference is 2 while the average total demand in one ordering period is 20.

Table 2. The impact of penalty cost.

Demand Distribution	Capacity/Mean Demand	Penalty/Holding Costs	Order-up-to-Levels
Poisson(5)	1.6	4.75	(0, 3, 6, 7)
Poisson(5)	1.6	12.25	(3, 6, 8, 8)
Poisson(5)	1.6	19.75	(5, 7, 9, 9)
Uniform(0, 1, . . . , 9)	1.67	4.75	(1, 4, 7, 8)
Uniform(0, 1, . . . , 9)	1.67	12.25	(5, 7, 9, 9)
Uniform(0, 1, . . . , 9)	1.67	19.75	(7, 8, 9, 9)
Binomial(0.5, 10)	1.6	4.75	(−1, 2, 5, 6)
Binomial(0.5, 10)	1.6	12.25	(1, 4, 6, 7)
Binomial(0.5, 10)	1.6	19.75	(2, 4, 6, 8)

Table 3. The impact of demand variability.

Demand Distribution	Coefficient of Variations	Order-up-to-Levels
Uniform(3, 4, 5, 6)	0.25	(−3, 1, 4, 6)
Uniform(2, 3, . . . , 7)	0.38	(−1, 2, 5, 7)
Uniform(1, 2, . . . , 8)	0.51	(3, 5, 7, 8)
Uniform(0, 1, . . . , 9)	0.64	(7, 8, 9, 9)

4.2. The Effect of Production Capacity

To explore the benefit of information sharing as a function of the production capacity, we illustrate in Figure 1 the percentage cost savings from information sharing with the optimal policy relative to no information sharing for five demand distributions.

For each demand distribution and each capacity level, we set the ratio of the penalty cost to holding cost to 4.75. Similar results can be obtained at other values of penalty over holding cost, and we will discuss the impact of penalty cost in Section 4.3.

Our computational study reveals that as production capacity increases, percentage cost savings increase. Indeed, percentage cost savings increase from about 10% to about 35% as capacity over mean demand increases from 1.2 to 3. This is quite intuitive, since as production capacity increases, the optimal policy would postpone production as much as possible and take advantage of all information available prior to the time production starts. For instance, in case of infinite capacity, it is optimal to wait until the last information period and produce to satisfy all demand realized so far plus an additional amount based on solving a newsboy problem (see Section 2.3). Similarly, if the production capacity is limited, then information is not very beneficial since the production quantity is mainly determined by capacity, not the realized demand.

Finally, from the fill-rate point of views, our computational study reveals that information sharing with the optimal policy and the no information sharing strategies have almost identical fill rates.

To explore the effectiveness of information sharing with the greedy policy, we provide in Figure 2 the percentage cost savings of information sharing with the optimal policy relative to information sharing with the greedy policy under similar conditions as above. The figure illustrates the following insights:

- Information sharing with the optimal policy reduces cost by at least 15% relative to information sharing with the greedy policy and the savings can be as much as 50–60%.
- When the production capacity is tightly constrained, the savings provided by

Table 4. Two ordering periods.

Demand Distribution	Capacity/Mean Demand	Penalty/Holding Costs	Order-up-to-Levels
Poisson(5)	1.6	4.75	(0, 3, 6, 7, 0, 3, 6, 8)
Poisson(5)	2	4.75	(−6, 1, 4, 7, −6, 1, 4, 8)
Uniform(0, 1, . . . , 9)	1.67	4.75	(1, 4, 7, 8, 1, 4, 7, 8)
Uniform(0, 1, . . . , 9)	2.11	4.75	(−5, 0, 5, 8, −5, 0, 5, 9)
Binomial(0.5, 10)	1.6	4.75	(−1, 2, 5, 6, −1, 2, 5, 8)
Binomial(0.5, 10)	2	4.75	(−7, −2, 3, 6, −8, −2, 3, 7)

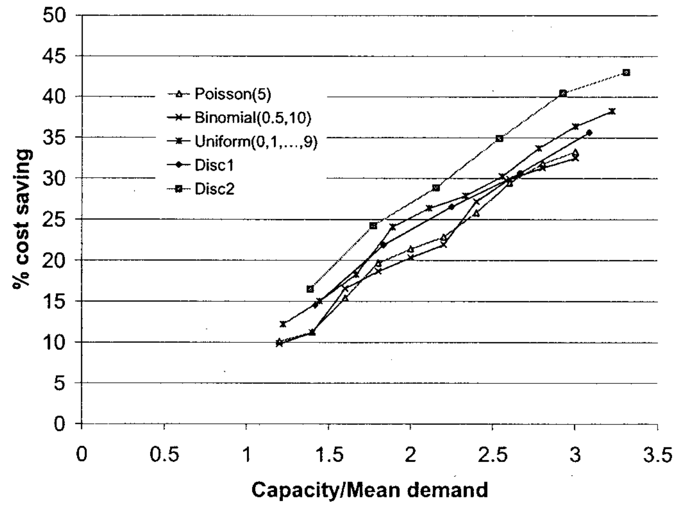

Figure 1. The impact of production capacity.

information sharing with the optimal policy is relatively high. This is because the greedy policy does not build safety stock until the last information period. In the last information period, if capacity is very tight, the greedy policy may not be able to build as much safety stock as needed, thus resulting in heavy penalty cost. On the other hand, the optimal policy can start building safety stock from the beginning of the ordering period by taking advantage of excessive capacity in all information periods.

- As capacity increases, the benefit from information sharing with the optimal policy relative to information sharing with the greedy policy decreases first and then

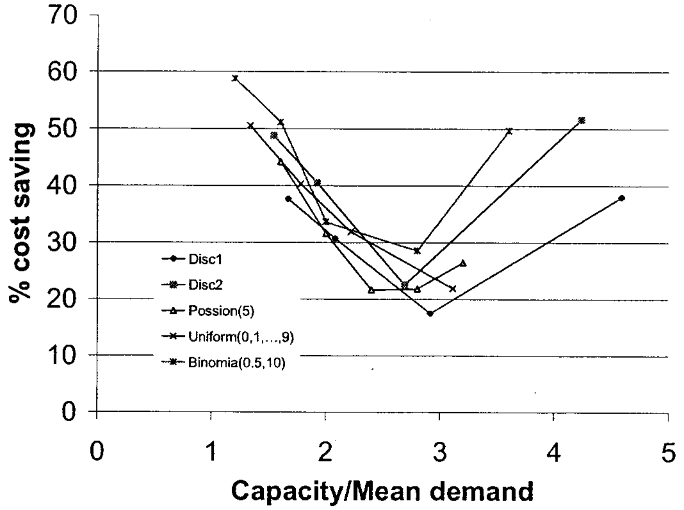

Figure 2. The benefits of using information optimally.

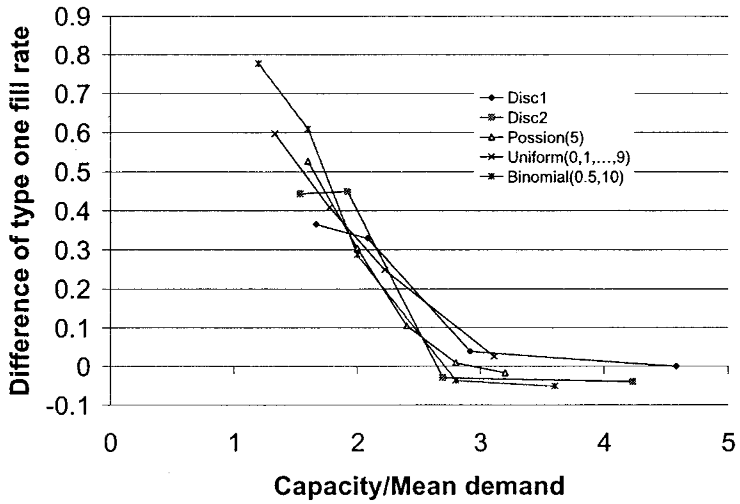

Figure 3. Fill rates: information sharing with the optimal policy vs. greedy policy.

increases again. This is true, since as capacity becomes very large relative to average demand, information sharing with the optimal policy will postpone production as much as possible, while information sharing with the greedy policy will build inventory starting from the beginning of ordering periods, thus resulting in heavy inventory holding cost.

Figure 3 shows the difference between type-one fill rates for information sharing with the optimal policy and information sharing with the greedy policy as a function of the production capacity for various demand distributions. The figure demonstrates that when capacity is relatively tight, the difference in the fill-rates is substantial. However, as capacity increases, the two strategies have almost identical fill-rates.

4.3. The Effect of Penalty Cost

To study the impact of penalty cost on the benefit of information sharing, we present in Figure 4 the percentage cost savings with information sharing relative to no information sharing as a function of the *ratio* of penalty cost to inventory holding cost at various capacity levels.

Demand distribution in one information period is assumed to be $Uniform(0, 1, \ldots, 9)$. As Gavirneni, Kapuscinski, and Tayur [4] points out, information sharing has limited value when the ratio is either very high or very low. For moderate ratios, Figure 4 illustrates that the percentage cost saving depends strongly on production capacity. Specifically, when capacity is tightly constrained (e.g., capacity/mean demand = 1.2), the percentage cost saving reaches its peak value at a smaller penalty/holding cost ratio than those when capacity is not tightly constrained (e.g., capacity/mean demand = 1.67, 2.11, 3). This is explained as follows: When capacity is tightly constrained, the total cost for both the no information sharing and information sharing strategies increases quite fast as penalty costs increase. Thus, the percentage saving reaches its peak value at a small ratio of penalty cost to inventory holding cost.

4.4. The Effect of the Number of Information Periods

To explore the benefits of information sharing as a function of the number of information periods in one ordering period, we present in Figure 5 the percentage cost savings with

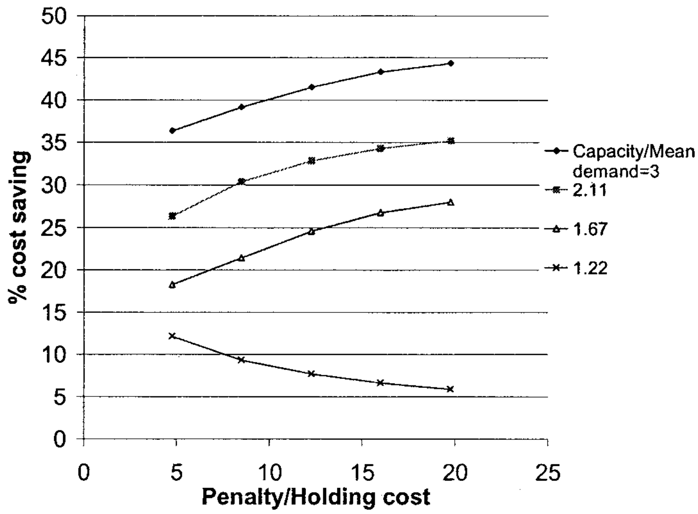

Figure 4. The impact of penalty cost.

information sharing relative to no information sharing for two production capacity levels. The number of information periods, N, was 2, 4, 6, and 8 while the length of the ordering period was assumed to be constant in all the models. The demand distribution during the entire ordering period is assumed to be $Poisson(\lambda)$ with $\lambda = 24$; hence demand in a single information period follows $Poisson(\lambda/N)$. Similarly, the total production capacity, and the inventory holding cost per item, in the entire ordering period are kept constant and are equally divided among the different information periods. The ratio of penalty to holding cost is 4.75.

Figure 5 illustrates the following insights:

- As the number of information periods increases, the percentage savings increase.

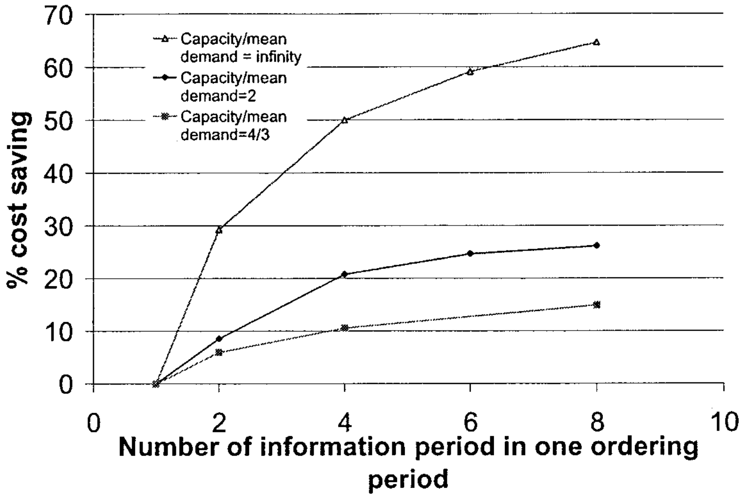

Figure 5. The impact of information sharing frequency.

- Most of the benefits from information sharing is achieved with a few information periods. That is, the marginal benefit is a decreasing function of the number of information periods. Specifically, the benefit achieved by increasing the number of information periods from 4 to 8 is relatively small.
- Define the maximum potential benefit from information sharing to be the percentage cost reduction when the manufacturer has unlimited capacity. A manufacturer with a production capacity twice as much as mean demand can achieve a substantial percentage of the maximum potential benefit, e.g., when the frequency of information sharing is 4, the manufacturer can obtain nearly 50% of the cost benefit that a manufacturer with unlimited production capacity can achieve.

4.5. Optimal Timing for Information Sharing

In this subsection we analyze the impact of the time(s) when information is shared on the manufacturer's total inventory and penalty costs.

4.5.1. Sharing Information Once in One Ordering Period

Figure 6 presents the manufacturer's total cost as a function of the time when information is shared, assuming that the manufacturer can only share information once in one ordering period. The figure provides the normalized manufacturer's cost as a function of a normalized time. That is, time is normalized and is measured from 0 to 1, while the cost is normalized by the cost of carrying one ordering period's total expected demand for one ordering period. Thus, 0 in the x coordinate implies that information is shared at the beginning of the ordering period, and 1 means that information is shared at the end of the ordering period and hence can not be utilized. Demand distribution is assumed to be $Poisson(24)$, and penalty over holding cost equals 4.

In Figure 7 we demonstrate the impact of the production capacity and penalty cost on the optimal timing of information sharing. These figures illustrate the following insights:

- As information sharing is delayed, the manufacturer's total cost first decreases and then increases sharply. The cost reaches its maximum when information is shared at the beginning or end of one ordering period.
- The optimal timing for information sharing is not in the middle of the ordering period for any combination of production capacity and penalty cost; rather, it is in the later half of the ordering period.
- Both the production capacity and penalty cost have minor impacts on the optimal timing of information sharing. For all combination of production capacity and penalty cost, the optimal timing is somewhere between 0.75 and 0.9 of the normalized time, and very close to 0.8 on average.
- When capacity is very large, it is appropriate to postpone the timing of information sharing to the last production opportunity in this ordering period; interestingly, this is also the right thing to do when capacity is tightly constrained, i.e., postpone the timing of information sharing until the last production opportunity. While the first observation (i.e., large capacity) is quite intuitive, the second one (tight capacity) is less intuitive. One possible explanation is that when capacity is very tight, the manufacturer needs to build as much inventory as she can until the last production opportunity, when she can review demand information and adjust production quantity.

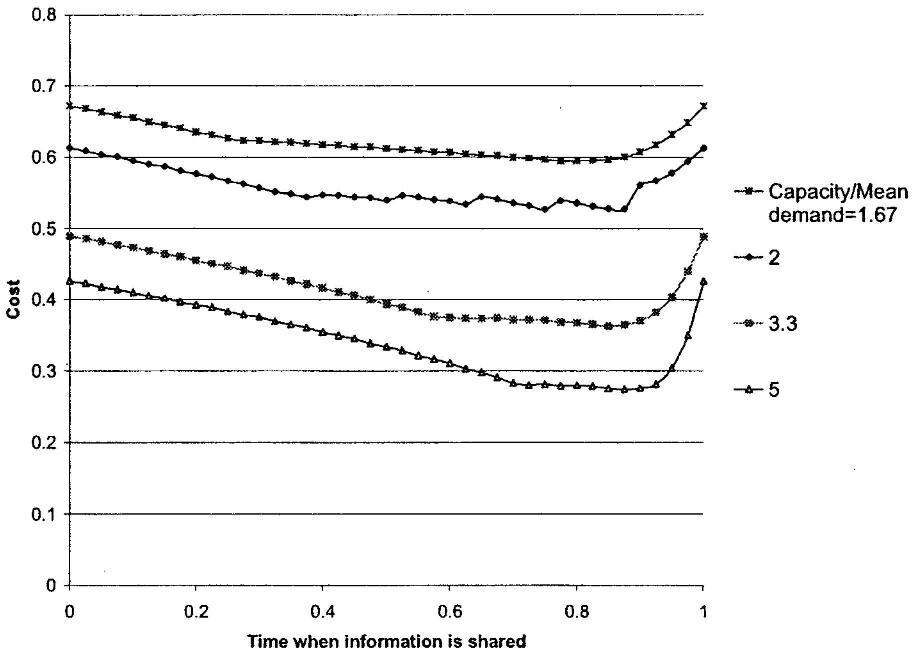

Figure 6. The impact of the timing of information sharing on manufacturer's cost.

4.5.2. Sharing Information Twice in One Ordering Period

Figures 8 and 9 present the manufacturer's total cost as a function of t_1/T and t_2/T, assuming that the manufacturer can share information at times t_1 and t_2 in a single ordering period. Similar to the previous section, the figure provides the normalized manufacturer's cost as a function of

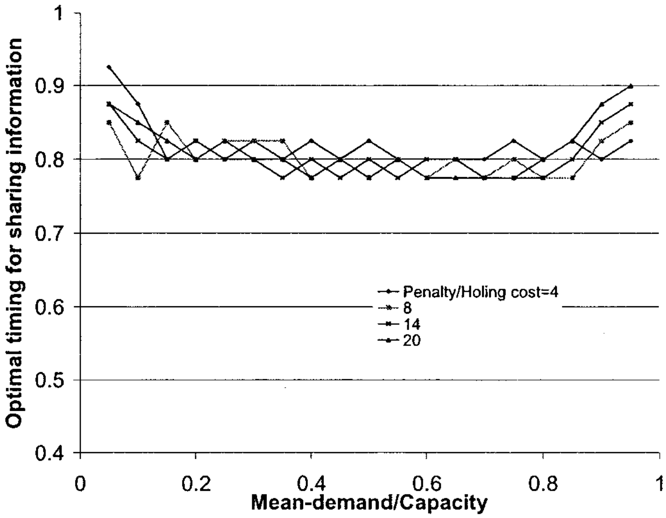

Figure 7. Optimal timing for information sharing with different penalty cost.

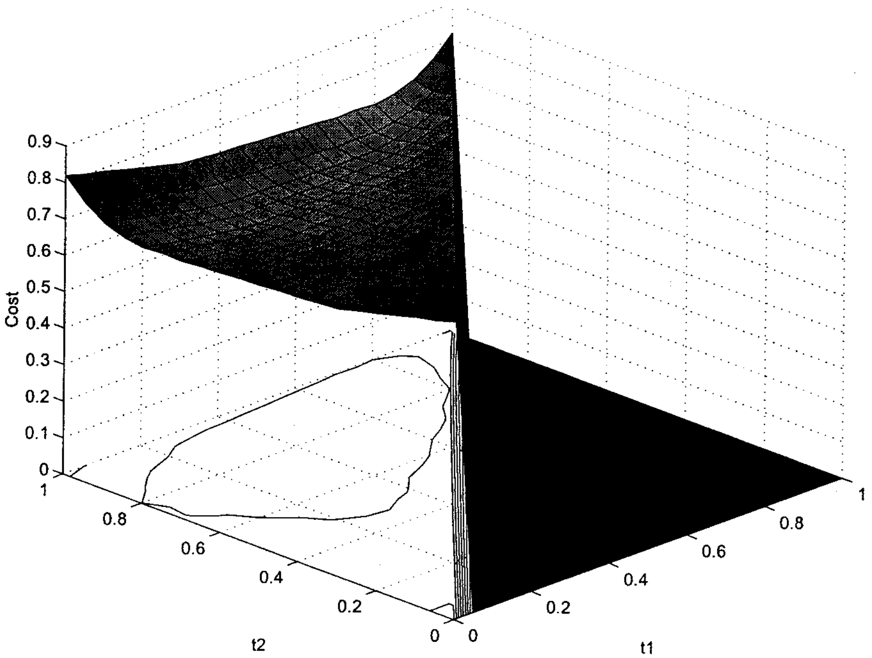

Figure 8. The impact of the timing of information sharing on manufacturer's performance.

a normalized time. Demand distribution is assumed to be $Poisson(10)$, penalty over holding cost equals 4, and the production capacity is twice the mean demand.

In Figure 10 we study the impact of the production capacity and penalty cost on the optimal timings of information sharing. The solid curves are the optimal timings for the second information sharing, and the dash curves are the optimal timings for the first information sharing. These figures illustrate the following insights:

- Sharing information twice in an ordering period helps the manufacturer achieve more benefits than sharing information once; e.g., the manufacturer's maximum cost saving increases from 14.6% (when information is shared only once, i.e., Fig. 9, $t_2 = 1$, $t_1 = 0.8$) to 20.7% (when it is shared twice, i.e., $t_2 = 0.8$, $t_1 \sim 0.3$).
- The optimal t_1 varies significantly as a function of the ratio of mean demand to production capacity, and it approaches the end of the ordering period when the available capacity is very large, i.e., when the ratio tends to zero. The ratio of penalty to inventory holding cost also has an impact on the optimal value of t_1 especially when the capacity utilization is neither very high nor very low.
- Unlike the optimal t_1, the optimal value of t_2 is not effected as much by either the ratio of mean demand to production capacity or the ratio of penalty to inventory holding cost.
- The optimal value of t_2 is unlikely to be in the first half of the ordering period for all combinations of parameters. On the other hand, the optimal value of t_1 can be either in the first or the second halves of the ordering period. Indeed, in our computational results the optimal value of t_1 varied from 0.22 to 0.82.

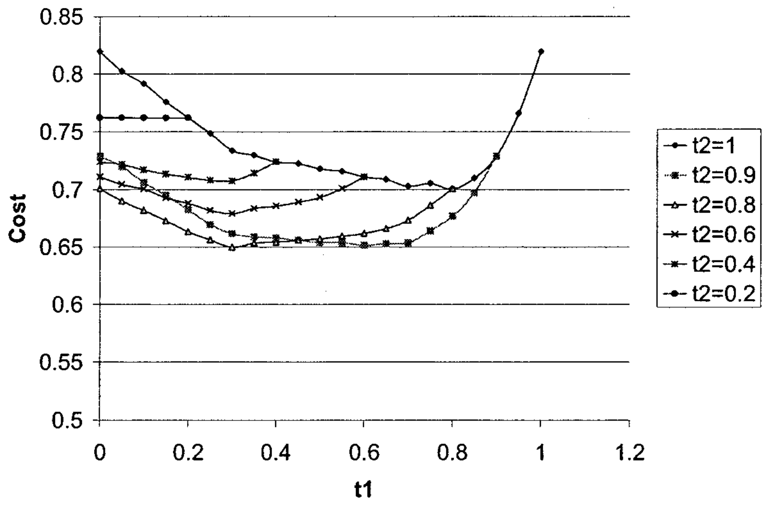

Figure 9. The impact of the timing of information sharing on manufacturer's performance.

5. CONCLUSION

In this paper, we consider a two-stage supply chain with a single retailer facing i.i.d demand and a single manufacturer with finite production capacity. In the model, the manufacturer receives demand information from the retailer even during time periods in which the retailer does not order. By analyzing the model in a finite time horizon, we study the value of information sharing for the manufacturer as well as how the manufacturer can utilize the shared demand information effectively.

As for the optimal inventory control policy under information sharing, we find that the following property holds under certain conditions: The optimal order-up-to-levels in any ordering period are nondecreasing as we move from the first information period to the last one.

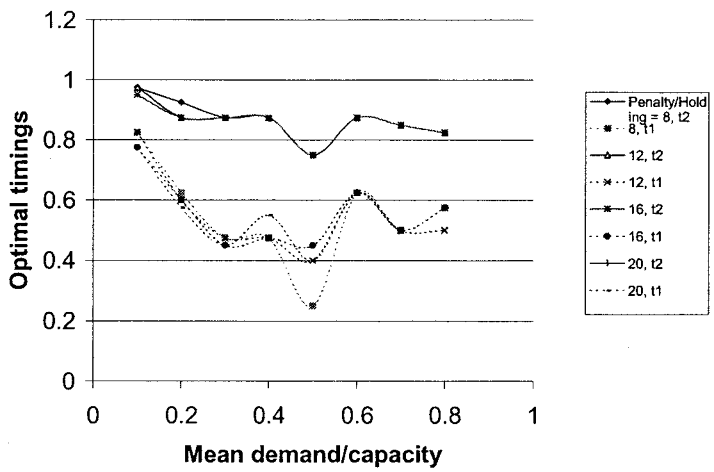

Figure 10. The impact of capacity and penalty cost on the optimal timings of information sharing.

We demonstrate, through an extensive computational study, the potential benefits of sharing demand information in terms of the manufacturer's cost and service level. For instance, if the manufacturer has excessive capacity, information sharing can be very beneficial. Indeed, the manufacturer can cut down on inventory costs while maintaining the same service level to the retailer by using information effectively. One interesting observation is that the manufacturer can realize most of the benefits from information sharing if the retailer shares demand information with the manufacturer only a few times in each ordering period.

If the retailer has only one opportunity to share information with the manufacturer, then the best time to share information is in the later half of the ordering period. Parameters such as the production capacity and penalty cost have limited impact on the optimal timing. If the retailer has two opportunities to share information with the manufacturer, the best timing of first information transferring is sensitive to the changes of capacity and penalty cost, while the optimal timing of the second information sharing is not.

It is appropriate to point out some limitations of our model. First, our focus in this model is exclusively on the impact of information sharing on the manufacturer but not the retailer. Second, the assumption that the retailer controls her inventory using an order-up-to policy with a constant order-up-to level, i.e., a level that does not change over time, implicitly implies our assumption that the retailer is not able to anticipate future shortages at the manufacturer. Of course, if the retailer is able to anticipate future shortages, then the retailer may inflate her orders so that she will receive the amount she needs. This type of behavior is not captured by our model.

APPENDIX A: PROOF OF LEMMA 1

The proof of parts (a) and (b) are identical to the one in Federgruen and Zipkin [7]. Here we focus on the proof of part (c).

Let us first consider the last ordering period in the planning horizon, and rewrite the dynamic programming formulation,

$$U'_n(x) = -(c + h_{n-1})x + V'_n(x),$$

$$V'_n(x) = \min_{x \le y \le x+C}\{J'_n(y)\}$$

$$J'_n(y) = \begin{cases} cy + h_{n-1}y + \beta U'_{n-1}(y), & n = 2, \dots, N, \\ cy + E(L(y, \sum D)) + \beta E(U'_{n-1}(y - \sum D)), & n = 1, \end{cases}$$

as

$$\begin{aligned} J'_n(y) &= cy + h_{n-1}y + \beta U'_{n-1}(y) \\ &= (1-\beta)cy + (h_{n-1} - \beta h_{n-2})y + \beta V'_{n-1}(y) \\ &= (1-\beta)cy + (h_{n-1} - \beta h_{n-2})y + \beta \min_{y \le y' \le y+C}\{J'_{n-1}(y')\}, \qquad n = 2, \dots, N, \\ J'_1(y) &= cy + E(L(y, \textstyle\sum D)). \end{aligned}$$

The last equation holds since the salvage cost $U_0(\cdot) \equiv V_0(\cdot) \equiv 0$. To simplify the notation, define

$$r'_n(y) = \begin{cases} (1-\beta)cy + (h_{n-1} - \beta h_{n-2})y, & 2 \le n \le N, \\ cy + E(L(y, \sum D)), & n = 1. \end{cases}$$

The minimal total cost for this ordering period $J'_N(y_N)$ can be obtained by plugging in the equation of $J'_n(y)$ recursively for $n = N - 1, \dots, 1$.

$$J'_N(y_N) = r'_N(y_N) + \beta \min_{y_N \le y_{N-1} \le y_N + C}\{r'_{N-1}(y_{N-1}) + \cdots + \beta \min_{y_3 \le y_2 \le y_3 + C}\{r'_2(y_2) + \beta \min_{y_2 \le y_1 \le y_2 + C}\{r'_1(y_1)\}\} \cdots\}. \tag{7}$$

Notice that for all $n = N, N - 1, \ldots, 1$, if $y + \Delta_n$ minimizes $r'_n(y')$ in interval $[y, y + C]$, then

$$|\min_{y \le y' \le y+C}\{r'_n(y')\} - r'_n(y)| = |r'_n(y + \Delta_n) - r'_n(y)| \propto |\Delta_n|,$$

since $r'_n(y)$ is at most proportional to a linear function of y. Because $\Delta_n \le C$, the absolute difference between (7) and $\sum_{n=1}^{N} \beta^{N-n} r'_n(y_N)$ is bounded by a finite number, which is independent of y. Finally, expanding $\sum_{n=1}^{N} \beta^{N-n} r'_n(y)$, we obtain $\sum_{n=1}^{N} \beta^{N-n} r'_n(y) = cy + h_{N-1}y + \beta^{N-1}E(L(y, \Sigma D))$. Hence, in order for $J_n(y) \to +\infty$ for all n when $y \to -\infty$, we need $\beta^{N-1}\pi > c + h_{N-1}$.

The proof for other ordering periods is similar. □

APPENDIX B: PROOF OF PART (c) OF LEMMA 3

The case when $y \to +\infty$ is straightforward. Thus, we only need to show that if $\beta^{N-1}\pi > c + h_{N-1}$, then $J_n(y) \to +\infty$ for all $n = N, N - 1, \ldots, 1$ when $y \to -\infty$.

For this purpose, let us consider the last ordering period in the planning horizon. From Eq. (4),

$$\begin{aligned} J_n(y) &= \varphi_n(y) + \beta E(U_{n-1}(y - D)) \\ &= r_n(y) + \beta E(V_{n-1}(y - D)) \\ &= r_n(y) + \beta E(\min_{y-D \le y' \le y-D+C}\{J_{n-1}(y')\}), \qquad 2, \ldots, N \\ J_1(y) &= r_1(y), \end{aligned}$$

where the salvage cost $U_0(\cdot) \equiv V_0(\cdot) \equiv 0$, and

$$r_n(y) = \begin{cases} (1 - \beta)cy + (h_{n-1} - \beta h_{n-2})y + \beta(c + h_{n-2})\mu, & 2 \le n \le N, \\ cy + E(L(y, D)), & n = 1. \end{cases}$$

The minimal total cost for this ordering period $J_N(y_N)$ can be obtained by plugging in the equation of $J_n(y)$ recursively for $n = N - 1, \ldots, 1$.

$$J_N(y_N) = r_N(y_N) + \beta E(\min_{y_N - D \le y_{N-1} \le y_N - D + C}\{r_{N-1}(y_{N-1}) + \cdots + \beta E(\min_{y_2 - D \le y_1 \le y_2 - D + C}\{r_1(y_1)\}) \cdots\}). \tag{8}$$

Similarly to Appendix A, we observe that for all $n = N, N - 1, \ldots, 1$,

$$|E(\min_{y-D \le y' \le y-D+C}\{r_n(y')\}) - r_n(y)| = |E(r_n(y - \Delta_n(D))) - r_n(y)|,$$

with $D - C \le \Delta_n(D) \le D$, and

$$\begin{aligned} |E(r_n(y - \Delta_n(D))) - r_n(y)| &= \left| \int_{-\infty}^{+\infty} [r_n(y - \Delta_n(D)) - r_n(y)] f_D(D)\, dD \right| \\ &\propto \left| \int_{-\infty}^{+\infty} \Delta_n(D) f_D(D)\, dD \right|, \end{aligned}$$

since $r_n(y)$ is at most proportional to a linear function of y. Because $|\int_{-\infty}^{+\infty} \Delta_n(D) f_D(D)\, dD| \le \max\{|\mu - C|, |\mu|\}$, the absolute difference between (8) and $\sum_{n=1}^{N} \beta^{N-n} r_n(y_N)$ is bounded by a finite number, which is independent of y. Finally, expanding $\sum_{n=1}^{N} \beta^{N-n} r_n(y_N)$ we obtain $\sum_{n=1}^{N} \beta^{N-n} r_n(y) = cy + h_{N-1}y + \beta^{N-1}E(L(y, D)) + \textit{constant}$. Hence, in order to have $J_n(y) \to +\infty$ for all n when $y \to -\infty$, we need $\beta^{N-1}\pi > c + h_{N-1}$.

The proof for other ordering periods is similar. □

ACKNOWLEDGMENTS

Research was supported in part by ONR Contracts N00014-95-1-0232 and N00014-01-1-0146, and by NSF Contracts DDM-9322828 and DMI-9732795.

REFERENCES

[1] Y. Aviv and A. Federgruen, The operational benefits of information sharing and vendor managed inventory (VMI) programs, Working Paper, Washington University, St. Louis, MO, 1998.
[2] G.P. Cachon and M. Fisher, Supply chain inventory management and the value of shared information, Management Sci 46 (2000), 1032–1050.
[3] A. Federgruen and P. Zipkin, An inventory model with limited production capacity and uncertain demands 2: The discounted-cost criterion, Math Oper Res 11 (1986), 208–215.
[4] S. Gavirneni, R. Kapuscinski, and S. Tayur, Value of information in capacitated supply chains, Management Sci 45 (1999), 16–24.
[5] G. Hadley and T.M. Whitin, Analysis of inventory systems. Prentice-Hall, Englewood Cliffs, NJ, 1963.
[6] D. Heyman and M. Sobel, Stochastic models in operations research, McGraw-Hill, New York, 1984, Vol. 2.
[7] H.L. Lee and S. Nahmias, "Single-product, single-location models," Handbooks in operations research and management science, North-Holland, Amsterdam, 1993, Vol. 4.
[8] D. Simchi-Levi, F. Kaminsky, and E. Simchi-Levi, Designing and managing the supply chain, Irwin/McGraw-Hill, Chicago, IL, 1999.
[9] T. Stein and J. Sweat, Killer supply chains, Inform Week (1998), http://www.informationweek.com/708/08iukil.htm.