

Supply Chain Inventory Management and the Value of Shared Information

Gérard P. Cachon • Marshall Fisher

*The Wharton School, The University of Pennsylvania, Philadelphia, Pennsylvania 19104
fisher@opim.wharton.upenn.edu*

In traditional supply chain inventory management, orders are the only information firms exchange, but information technology now allows firms to share demand and inventory data quickly and inexpensively. We study the value of sharing these data in a model with one supplier, N identical retailers, and stationary stochastic consumer demand. There are inventory holding costs and back-order penalty costs. We compare a traditional information policy that does not use shared information with a full information policy that does exploit shared information. In a numerical study we find that supply chain costs are 2.2% lower on average with the full information policy than with the traditional information policy, and the maximum difference is 12.1%. We also develop a simulation-based lower bound over all feasible policies. The cost difference between the traditional information policy and the lower bound is an upper bound on the value of information sharing: In the same study, that difference is 3.4% on average, and no more than 13.8%. We contrast the value of information sharing with two other benefits of information technology, faster and cheaper order processing, which lead to shorter lead times and smaller batch sizes, respectively. In our sample, cutting lead times nearly in half reduces costs by 21% on average, and cutting batches in half reduces costs by 22% on average. For the settings we study, we conclude that implementing information technology to accelerate and smooth the physical flow of goods through a supply chain is significantly more valuable than using information technology to expand the flow of information.

(Supply Chain; Multi-Echelon Inventory Management; Periodic Review Policies; Electronic Data Interchange)

1. Introduction

Information technology has had a substantial impact on supply chains. Scanners collect sales data at the point-of-sale, and electronic data interchange (EDI) allows these data to be shared immediately with all stages of the supply chain. The application of these technologies, especially in the grocery industry, has substantially lowered the time and cost to process an order, leading to impressive improvements in supply chain performance (see Cachon and Fisher 1997, Clark and Hammond 1997, Kurt Salmon Associates 1993).

As a result of these success stories, there is now a

general belief within industry that capturing and sharing real-time demand information is the key to improved supply chain performance. The purpose of this research is to test this belief by rigorously measuring the value of information sharing and comparing this value to two other sources of supply chain improvement: reducing lead times and increasing delivery frequency by reducing shipment batch sizes. Note that the same information technology that facilitates information sharing also contributes to the reduction of lead times and shipment frequency by reducing the time and cost to process orders. Thus, the

question addressed in here is not whether information technology improves supply chain performance, but how. Specifically, does the primary gain come from sharing information or from allowing products to flow more quickly and evenly in the supply chain?

We address this question within the context of a supply chain with one supplier and N identical retailers that face stationary stochastic consumer demand with a known distribution. There are fixed transportation times between locations, and shipment quantities equal a multiple of a base batch quantity. There are holding costs at all levels and back-order penalty costs at the lowest level. This model provides a reasonable representation of supply chains selling an established product under constant pricing conditions.

We consider two levels of information sharing. With *traditional information* sharing the supplier only observes the retailers' orders. With *full information* sharing the supplier has immediate access to the retailers' inventory data. We develop an inventory policy for each information sharing level. Reorder point policies are used with traditional information sharing. The retailers also use reorder point policies with full information, but the supplier does not. Instead, the supplier uses its additional information to better allocate inventory among the retailers and to improve its order decisions (i.e., to better time its own replenishments).

The difference between supply chain costs under traditional and full information is one measure of the value of shared information. However, there may exist even better policies for either information level, that is, optimal policies are unknown for each level. To account for this possible bias, we develop a simulation-based lower bound over all feasible policies, no matter what the level of information sharing is. The cost difference between traditional information and the lower bound is the maximum value of shared information.

In a numerical study with a wide range of parameter values we find that information sharing reduces supply chain costs by 2.2% on average, and the gap between traditional information policy cost and the lower bound is 3.4% on average. Cutting lead time by

nearly half (from five to three periods) reduces costs by 21% on average, and cutting batch size in half reduces supply chain costs by 22%. We recognize that this comparison is meaningful only if those lead time and batch size reductions can be reasonably expected from the implementation of information technology. In fact, we did observe comparable reductions at Campbell Soup Company when it implemented information technology to improve its supply chain.¹ There has also been other documentation on the impact of information technology in the grocery industry: Barilla, the world's largest pasta producer, reduced its lead time from over one week to two days (Harvard Business School case 9-694-046); and H.E.B., a large grocery chain based in Texas, eliminated 6 to 10 days from its lead time (Harvard Business School case 9-195-125). We conclude that while information sharing does reduce costs, simply flowing goods through the supply chain more quickly and more evenly produces an order of magnitude greater improvement.

The next section reviews the related literature. Section 3 outlines the model, and §4 describes how to select inventory policies with traditional information sharing. Section 5 details a lower bound over all feasible policies. Section 6 develops the full information inventory policy. Section 7 details a numerical study, §8 discusses our results, and §9 concludes.

2. Literature Review

The following papers show how sharing demand and inventory data can improve the supplier's order quantity decisions in models with known and stationary retailer demand: Bourland et al. (1996), Chen (1998),

¹ Campbell Soup Company gave us ordering data from several retailers before and after the implementation of information technology. In the "before" data we saw that retailers often purchased in multiple pallet quantities because, according to our contacts at Campbell Soup Company, they did not want to bother with the hassle of placing orders frequently. In the "after" data it is clear that each products' minimum batch size was no greater than one pallet, and in some cases Campbell Soup Company was willing to deliver in half-pallet increments. In addition, the lead time for deliveries to the retailers was reduced from about one week to two to three days, primarily resulting from the reduction in the order processing time.

Gavirneni et al. (1999), and Aviv and Federgruen (1998). Lee et al. (2000) use shared information to improve the supplier's order quantity decisions in a serial system with a known autoregressive demand process. Liljenberg (1996) studies how to use shared information to improve the supplier's allocation of inventory among the retailers. In our model shared information is exploited for both uses: better supplier replenishments and better allocations to the retailers.

We focus on sharing demand and inventory data, but there are other data that can be shared in a supply chain. Gavirneni et al. (1999) measure the benefit of sharing the parameters of the retailer's ordering policy with the supplier. Aviv (1998) explores the benefits of sharing forecasts for future demand.

In our model, as in the other studies mentioned, it is assumed that information is always shared truthfully. Cachon and Lariviere (1997) study forecast sharing when the forecast provider has an incentive to provide an overly optimistic forecast of demand.

Both Lee et al. (2000) and Gavirneni et al. (1999) assume there exists a perfectly reliable exogenous source of inventory; information sharing has no impact on the retailer because its orders are always received in full after a fixed number of periods. In the other papers, as in our model, the supplier is the only source of inventory. Therefore, information sharing may impact the retailers by changing the supplier's order quantities or allocations.

Gavirneni et al. (1999) and Aviv and Federgruen (1998) allow for limited supplier capacity, whereas capacity is unrestricted in our model and in the other papers.

The reported benefits of information sharing vary considerably. Liljenberg (1996) finds that better allocation lowers supply chain costs by 0% to 3.9%. Chen (1998) finds that supply chain costs are lowered up to 9%, and on average by 1.8%. Aviv and Federgruen (1998) report benefits of 0%–5%. In contrast, Lee et al. (2000) find that information sharing lowered supply chain costs by about 23% in their scenario with the highest demand nonstationarity. However, Graves (1999) studies a similar model, with the exception that there is no outside inventory source, and concludes that information sharing provides no benefit to the

supply chain. Gavirneni et al. (1999) report that sharing the retailer's demand data reduced the *supplier's* cost by 1%–35%.² The impact on the *supply chain's* cost would be lower because information sharing in their model has no impact on the retailer's costs.

There is other research related to our work. Lee et al. (1997) find that sharing information reduces the supplier's demand variance, which should benefit the supply chain, but they do not quantitatively measure this benefit. There are many studies that investigate a supply chain model with one supplier, N retailers, stochastic consumer demand, and batch ordering. Some of them assume traditional information (e.g., Axsäter 1993, Cachon 1995, Chen and Samroengraja 1996, Lee and Moinzadeh 1986, Svoronos and Zipkin 1988), while others assume full information (e.g., Chen and Zheng 1997, Graves 1996, McGavin et al. 1993). Because of different assumptions and test problems, it is not possible to meaningfully compare supply chain costs across those two sets of studies. Several researchers study allocation rules, but none addresses the issue of information sharing (see Cachon 1995, Chen and Samroengraja 1996, and Graves 1996). Anand and Mendelson (1997) study a one-period model in which retailers possess some local information that cannot be shared with either a central agent or other retailers. In our full information model all relevant information can be shared with the central agent (i.e., the supplier).

Chen and Zheng (1994) develop a lower bound over all feasible policies for a multiple retailer model. They show that a full information policy is reasonably close to optimal, but they do not compare this policy with a traditional information policy. Graves (1996) also shows that his full information policy is close to optimal.

3. Model

Firm 0 is the supplier and firms $[1, N]$ are identical retailers. The supplier replenishes its inventory from a perfectly reliable single source, that is, the supplier's orders are always received after a constant lead time.

² Those data are taken from the scenarios with the highest supplier capacity. They found that the value of information sharing declined as the supplier's capacity decreased.

The supplier is the retailers' only source of inventory, so there is no diversion of stock among the retailers and stockouts at the supplier cause replenishment delays for the retailers. Inventory is reviewed periodically, and within each period the following sequence of events occurs: (1) retailers order, (2) the supplier orders, (3) inventory shipments are received and released, (4) consumer demand occurs, and (5) inventory holding and backorder costs are charged.

The following variables are defined before demand in period t (between Events 3 and 4) for each firm $i \in [0, N]$: $I_i(t)$, on-hand inventory; $B_i(t)$, outstanding back orders; $O_i(t)$, on-order inventory, but not yet shipped; $IT_i(t)$, in-transit inventory; and $IP_i(t) = I_i(t) - B_i(t) + IT_i(t)$, inventory position.³ The supplier's on-order inventory is always zero because its source always ships inventory immediately, but a retailer's on-order inventory is positive when the supplier is unable to fill the retailer's order completely. Variables without a period designation apply to period t , e.g., I_i is retailer i 's inventory in period t before consumer demand.

A superscript "e" on a variable indicates that it is measured at the end of the period, and a superscript "b" indicates that it is measured at the beginning of the period. (Naturally, $IP_i^e(t) = IP_i^b(t + 1)$.) Dropping the subscript from a variable denotes the vector of that variable across all firms, e.g., $IP = \{IP_{0^r}, \dots, IP_{N^r}\}$. Since retailers are identical, it is convenient to associate a variable with a generic retailer: A subscript "r" indicates a variable applies to any retailer, and a subscript "s" identifies the supplier.

Define D_r^τ as demand at one retailer over τ consecutive periods. Demand is discrete, independent, and identically distributed across retailers and across time. We assume $D_r^1 \in [0, \bar{d}]$, \bar{d} is a finite integer, and $\Pr(D_r^1 = 1) > 0$. Let $\mu_r = E[D_r^1]$. Demands not filled immediately from stock are backordered and eventually filled.

There are holding and back-order costs in each period: $h_s > 0$ per unit of stock at the supplier or en

route to the retailers, $h_s + h_r$ per unit of stock at the retailers, where $h_r \geq 0$, and p_r per back order at each retailer.

Each shipment to a retailer equals an integer number of *batches*, where a batch is Q_r units. Each supplier order is an integer multiple of $Q_s Q_r$ units. Because all of the supplier's shipments to the retailers equal an integer number of batches, all supplier variables are measured in batches. When the supplier submits an order in period t , it receives the entire order in period $t + L_s$. A batch shipped to a retailer in period t is received in period $t + L_r$. The supplier will ship part of a multiple batch retailer order.

4. Traditional Information Policies

This section defines the inventory policies the firms use under traditional information. We then demonstrate how to evaluate expected supply chain cost for a given policy. A search over the set of reasonable policies (defined later) yields the optimal policy.

4.1. Inventory Policies

The retailers implement a (R_r, nQ_r) reorder point policy: When $R_r \geq IP_r^b + O_r^b$, a retailer orders the smallest integer multiple of Q_r units to ensure that $R_r < IP_r + O_r$. The supplier implements an analogous (R_s, nQ_s) reorder point policy. Reorder point policies are probably not optimal in most cases, but the optimal policies under traditional information are not known. Nevertheless, reorder point policies are simple to implement and intuitively reasonable. In fact, they are optimal in a serial supply chain with batch ordering (Chen 1997).

It remains to define the supplier's policy for allocating inventory among the retailers when the supplier's inventory is insufficient to cover the retailers' total orders in a period. We assume the supplier implements *batch priority* allocation. With this scheme the supplier assigns in every period a priority to each batch ordered in that period. Suppose retailer i orders b batches in a period. Then, the first batch in the order is assigned priority b , the second batch is assigned priority $b - 1$, and so forth. All batches ordered within a period are placed in a shipment queue, and they enter in decreasing priority order. If multiple

³ Some authors define inventory position to include on-order inventory. That definition is less cumbersome when working with traditional ordering policies but more cumbersome when working with the lower bound.

batches have the same priority, a fair coin toss determines the sequence in which they enter the shipment queue. Actual deliveries to the retailers are based on the shipment queue on a first-in-first-out basis. Therefore, all batches ordered before period t are shipped ahead of batches ordered in period t (but possibly in the same period). Under batch priority the retailer ordering the most batches within a period is presumed to have the highest need, and therefore the first batch in its order is assigned the highest priority for the period. Of course, priorities only matter when inventory is insufficient to cover the retailers' orders.

Even though batch priority is a reasonable allocation rule given traditional information, there may exist better allocation schemes. For example, because the supplier observes the time since each retailer's last order, the supplier might use those data to better determine the retailers' priorities. However, we believe that it is computationally intractable to evaluate expected costs under that allocation scheme. Furthermore, the difference between our traditional information policy and the lower bound provides a conservative estimate of the value of shared information.

4.2. Evaluating Expected Costs

Let $E[C]$ be the expected per period supply chain holding and back-order costs for a given pair of reorder points,

$$E[C] = h_s Q_r E[I_s^e] + N(h_s E[IT_r^e] + (h_s + h_r) E[I_r^e] + p_r E[B_r^e]),$$

where all of the above expectations depend on the chosen reorder points. (They are omitted for notational clarity.) The equations we provide evaluate $E[C]$ exactly when $R_s \geq -1$, otherwise those same equations evaluate $E[C]$ approximately. We begin with several straightforward results. We then turn to the evaluation of the supplier's shipping delay, i.e., the number of periods the supplier delays shipping a batch in a retailer's order. Evaluation of that distribution function is the main challenge in the evaluation of $E[C]$.

From Little's Law, $E[IT_r^e] = \mu_r L_r$ and $E[I_r^e] = \mu_r E[S]$, where $E[S]$ is the expected number of periods in which a unit is charged holding costs at a retailer (i.e.,

its expected sojourn in inventory). Also using Little's Law and observing that with reorder point policies $IP_s + O_s$ and $IP_r + O_r$ are uniformly distributed on the intervals $[R_s + 1, R_s + Q_s]$ and $[R_r + 1, R_r + Q_r]$, respectively, it can be shown that

$$E[B_r^e] = E[I_r^e] + \mu_r (E[U] + L_r + 1) - \left(R_r + \frac{Q_r + 1}{2} \right),$$

$$E[I_s^e] = R_s + \frac{Q_s + 1}{2} + \frac{N\mu_r}{Q_r} (E[U] - L_s - 1),$$

where $E[U]$ is the expected number of periods the supplier delays shipping a batch to a retailer. (See Axsäter 1997 for additional details.)

It remains to evaluate $E[U]$ and $E[S]$. Some additional random variables are needed. When $R_r \geq IP_r^b + O_r^b$, define $V_i(t) = R_i - IP_i^b(t) - O_i^b(t)$. Call V_i the overshoot. From Cachon (1995),

$$\Pr(V_i = v) = \frac{\Pr(D_r^1 \leq Q_r + v) - \Pr(D_r^1 \leq v)}{\sum_{c=0}^{Q_r-1} (1 - \Pr(D_r^1 \leq c))}.$$

Let \bar{v} be the maximum overshoot, $\bar{v} = \bar{d} - 1$. Let $b_r(v)$ be the number of batches a retailer orders because of an overshoot v ,

$$b_r(v) = \left\lfloor \frac{v}{Q_r} \right\rfloor + 1.$$

Define U_{bv} as the number of periods the supplier delays shipping the b th batch in a retailer order triggered by an overshoot v . From Cachon (1995),

$$E[U] = \frac{\sum_{b=1}^{b_r(v)} \sum_{v=0}^{\bar{v}} E[U_{bv}] \Pr(V_r = v)}{\sum_{v=0}^{\bar{v}} \Pr(V_r = v) b_r(v)}.$$

Taking a weighted average over all units in the supply chain yields

$$E[S] = \frac{1}{Q_r} \frac{\sum_{b=1}^{b_r(v)} \sum_{c=1}^{Q_r} \sum_{v=0}^{\bar{v}} \sum_{u=0}^{\bar{u}} E[S_{bcvu}] \Pr(V_r = v) \Pr(U_{bv} = u)}{\sum_{v=0}^{\bar{v}} \Pr(V_r = v) b_r(v)}, \quad (1)$$

where \bar{u} is the maximum shipping delay and S_{bcvu} is the inventory sojourn of the c th unit in the b th batch of

an order that is triggered by an overshoot v and experiences a shipping delay of u periods. (When $R_s \geq -1$, $\bar{u} = L_s - 1$, otherwise $\bar{u} \geq L_s - 1$.) From Cachon (1995),

$$E[S_{bcvu}] = \sum_{\tau=u+L_r+1}^{\infty} \Pr(D_r^\tau \leq R_r - v + (b-1)Q_r + c - 1). \quad (2)$$

When $R_s \geq -1$, (1) and (2) are exact, otherwise they are approximations: When $R_s < -1$, a retailer's lead time demand is not independent of the lead time.

Now the only missing piece is the distribution function of U_{bv} . That distribution function is evaluated next.

4.3. The Supplier's Delay Distribution

Suppose in period t retailer j experiences an overshoot v . Thus, U_{bv} is the number of periods the supplier delays shipping the b th batch in retailer j 's order. For brevity, let *batch* b refer to that batch. Because the supplier implements a reorder point policy, Cachon (1995) shows that

$$\Pr(U_{bv} \leq u) = \frac{1}{Q_s} \sum_{q=1}^{Q_s} \Pr(U_{bvq} \leq u),$$

where

$$\Pr(U_{bvq} \leq u) = \begin{cases} \Pr(XB_v^{L_s-u} \leq R_s + q - b) & R_s + q - b \geq 0, \quad 0 \leq u \leq L_s \\ 1 & R_s + q - b \geq 0, \quad u \geq L_s + 1 \\ \Pr(XF_v^{u-L_s-1} > -1 - b_r(v) - (R_s + q - b)) & R_s + q - b < 0, \quad u \geq L_s + 1. \end{cases}$$

In the above, XB_v^τ is the number of batches the retailers order over periods $[t - \tau, t]$, excluding batches in period t shipped *after* batch b and all batches retailer j orders in period t . Similarly, XF_v^τ is the number of batches the retailers order over periods $[t, t + \tau]$, excluding batches in period t shipped *before* batch b and all batches retailer j orders in period t . Think of XB_v^τ and XF_v^τ as the supplier's "time backward" and "time forward" demand processes relative to retailer j 's period t order. Note that the above results do not depend on the allocation scheme the supplier implements. The evaluations of XB_v^τ and XF_v^τ do depend on the allocation scheme.

XB_v^τ and XF_v^τ are the summation of retailer j 's order process and the order processes of the other $N - 1$ retailers. Refer to those other $N - 1$ retailers as the "non- j " retailers. Hence, we begin with retailer j 's order process and then consider the order process of a non- j retailer. Finally, these processes are combined to yield XB_v^τ and XF_v^τ .

4.3.1. Retailer j 's Order Process. Define JF_v^τ as the number of batches retailer j orders over periods $[t$

$+ 1, t + \tau]$, and JB_v^τ as the number of batches retailer j orders over periods $[t - \tau, t - 1]$. Define the following random variable for any retailer i ,

$$W_i(t) = R_r + Q_r - IP_i(t) - O_i(t).$$

Given that every Q_r demand triggers a retailer order, the (W_i) demand before the start of period t triggers retailer i 's last order in period t or earlier. Also, the $(Q_r - W_i)$ th subsequent demand after the start of period t triggers retailer i 's first order to occur after period t . In general, let $Y_r^\tau(W_i(t))$ be the number of batches retailer i orders over periods $[t + 1, t + \tau]$,

$$Y_r^\tau(x) = \left\lfloor \frac{(x + D_r^\tau)^+}{Q_r} \right\rfloor.$$

Let w be the realization of $W_j(t)$,

$$w = v - \left\lfloor \frac{v}{Q_r} \right\rfloor Q_r.$$

It follows that $JF_v^\tau = Y_r^\tau(w)$.

Let $W_{jv}(t - 1)$ be $W_j(t - 1)$ conditional on retailer j 's period t overshoot. From Bayes theorem,

$$\Pr(W_{jv}(t - 1) = w) = \frac{\Pr(D_r^1 = v + Q_r - w)}{\Pr(V_j = v)}.$$

Note that the $(W_{jv}(t - 1) + 1)$ previous demand triggered the last batch ordered by retailer j in period $t - 1$ or earlier. Hence, $JB_v^\tau = Y_r^\tau(Q_r - 1 - W_{jv}(t - 1))$.

4.3.2. The Non- j Retailers' Order Process. Consider retailer k , where $k \neq j$. Define YB^τ as the number of batches retailer k orders over periods $[t - \tau, t]$, excluding any batches shipped *after* batch b . Define YF^τ as the number of batches retailer k orders over periods $[t, t + \tau]$, excluding any batches shipped *before* batch b . Let π equal batch b 's priority.

Begin with retailer k 's period t order process. Retailer k orders $Y_r^1(W_k(t - 1))$ batches in period t . Because $W_k(t - 1)$ is independent of V_j , $W_k(t - 1)$ is uniformly distributed on the interval $[0, Q_r - 1]$. According to batch priority, all batches retailer k orders in period t with higher than b priority are shipped before batch b . (Recall that batch b 's priority is b .) Any batch with priority b is shipped before batch b with a 0.50 probability. Hence, let $YB^0(W_k(t - 1))$ be YB^0 conditional on $W_k(t - 1)$,

$$YB^0(W_k(t - 1)) = 0.5(Y_r^1(W_k(t - 1)) - \pi)^+ + 0.5(Y_r^1(W_k(t - 1)) - \pi + 1)^+.$$

For a given $W_k(t - 1)$, retailer k 's order process in period t is independent of its order process in periods $[t - \tau, t - 1]$. Hence,

$$\Pr(YB^\tau \leq y) = \frac{1}{Q_r} \sum_{w=0}^{Q_r-1} \sum_{x=0}^y \Pr(Y_r^\tau(Q_r - 1 - w) = x) \times \Pr(YB^0(w) \leq y - x).$$

Now consider YF^τ . All batches with a priority lower than b are certainly shipped after batch b , and a batch with priority b is shipped after batch b with 0.50 probability. Let $YF^0(W_k)$ be YF^0 conditional on W_k . Because $Y_r^1(Q_r - 1 - W_k)$ is the number of batches retailer k orders in period t ,

$$YF^0(W_k) = 0.5 \min(b, Y_r^1(Q_r - 1 - W_k)) + 0.5 \min(b - 1, Y_r^1(Q_r - 1 - W_k)).$$

W_k is independent of V_j , so W_k is also uniformly distributed on the interval $[0, Q_r - 1]$. Given W_k , retailer k 's order process in period t is independent of its order process in periods $[t + 1, t + \tau]$. Hence,

$$\Pr(YF^\tau \leq y) = \frac{1}{Q_r} \sum_{w=0}^{Q_r-1} \sum_{x=0}^y \Pr(Y_r^\tau(w) = x) \times \Pr(YF^0(w) \leq y - x).$$

4.3.3. The Order Process of N Retailers. Define NB^τ and NF^τ as the $(N - 1)$ th fold convolutions of YB^τ and YF^τ , respectively;

$$XB_{vj}^\tau = JB_v^\tau + NB^\tau,$$

$$XF_{vj}^\tau = JF_v^\tau + NF^\tau.$$

Because the ordering processes of retailers are independent, the above are simple convolutions.

4.4. Finding Optimal Reorder Point Policies

For a fixed R_s , $E[C]$ is convex in R_r , so it is straightforward to search for the optimal R_r given R_s . However, $E[C]$ is not necessarily jointly convex in R_s and R_r . Hence, finding the optimal reorder points requires a search over feasible R_s values. Fortunately, it is possible to constrain the search region. There is no need to consider $R_s < -Q_s$, because then increasing R_s provides the retailers with a better lead time without increasing the supplier's inventory above zero. For a sufficiently large R_s there is no need to consider an even larger R_s because then the supplier (almost) always fills the retailers' orders immediately; any further increase in R_s would increase the supplier's inventory without reducing costs at the retailers. (See Cachon 1995 or Axsäter 1997 for additional details.)

5. Lower Bound

This section develops for this model a lower bound for supply chain costs that is independent of the level of information sharing. The first step divides the supply chain's costs into two components. The second step

evaluates a lower bound for each component. The final step sums the components' lower bounds to yield a lower bound for the supply chain.

5.1. A Division of Supply Chain Costs

Actual supply chain holding and back-order costs in period t are

$$C = h_s Q_r I_s^e + \sum_{i=1}^N [h_s I T_i^e + (h_s + h_r) I_i^e + p_r B_i^e]. \quad (3)$$

Define the supplier's *echelon inventory level*, IL_s , as the amount of inventory (in units) at the supplier or lower in the system minus total consumer backorders (before consumer demand in period t):

$$IL_s^e = Q_r I_s^e + \sum_{i=1}^N [I T_i^e + I_i^e - B_i^e].$$

IL_s is measured in units (i.e., not in batches, as are the other supplier variables) because the supplier's echelon inventory level is not necessarily an integer multiple of Q_r units. Rewrite (3) as

$$C = h_s IL_s^e + \sum_{i=1}^N [h_r I_i^e + (h_s + p_r) B_i^e]. \quad (4)$$

Each period, the supply chain must decide an inventory quantity to order from the supplier's source and quantities to ship to the retailers. These shipment release decisions are made before the realization of demand, yet the actual costs, (4), are incurred after demand. To link the shipment release decisions to actual costs, we construct a system of charges that are incurred before demand and equal actual costs in expectation. (We deliberately use "charges" to refer to our system of fees and use "costs" to refer to actual costs.)

In our system of charges, in period t retailer i is charged $G_r(IP_i, R^*)$ and the supplier is charged $G_s(IL_s, IP, R^*)$, where

$$\begin{aligned} \bar{G}_r(y) &= E[h_r(y - D_r^{L_r+1})^+ \\ &\quad + (h_s + p_r)(y - D_r^{L_r+1})^-], \end{aligned}$$

$$G_p(y, R) = \begin{cases} \bar{G}_r(y) - \bar{G}_r\left(y + Q_r\left(\left\lfloor \frac{R-y}{Q_r} \right\rfloor + 1\right)\right) & y \leq R \\ 0 & \text{otherwise,} \end{cases}$$

$$G_r(y, R) = \bar{G}_r(y) - G_p(y, R),$$

$$G_s(y, IP, R) = h_s(y - N\mu_r) + \sum_{i=1}^N G_p(IP_i, R),$$

and R^* is the largest integer, such that $\bar{G}_r(R^* + Q_r) \leq \bar{G}_r(R^*)$.

We now confirm that the expected sum of our charges equals expected actual costs. Because $G_p(IP_i, R^*)$ is merely a transfer payment between the supplier and the retailers,

$$\begin{aligned} G_s(IL_s, IP, R^*) + \sum_{i=1}^N G_r(IP_i, R^*) &= h_s(IL_s - N\mu_r) \\ &\quad + \sum_{i=1}^N E[h_r I_i^e(t + L_r) + (h_s + p_r) B_i^e(t + L_r)]. \end{aligned}$$

Over an infinite horizon

$$\begin{aligned} E[h_r I_i^e(t + L_r) + (h_s + p_r) B_i^e(t + L_r)] \\ = E[h_r I_i^e + (h_s + p_r) B_i^e], \end{aligned}$$

and within any period $E[h_s IL_s^e] = E[h_s(IL_s - N\mu_r)]$. So

$$E[C] = E[G_s(IL_s, IP, R^*)] + \sum_{i=1}^N E[G_r(IP_i, R^*)]. \quad (5)$$

Thus, actual costs have been divided into two components: The first component is the supplier's charges and the second component is the retailers' charges.

5.2. Evaluation of the Lower Bound

We seek a lower bound for both the supplier's charges and the retailers' charges. Begin with the retailers. Let $C_r = \min E[\bar{G}_r(IP_i)]$. Since $\bar{G}_r(y)$ is convex and all retailer shipments must equal a multiple of Q_r units, a reorder point policy minimizes expected costs. In fact, R^* is the optimal reorder point: If $IP_i \geq R^* + 1$,

shipping another batch to the retailer raises its costs; if $IP_i \leq R^*$, shipping another batch to the retailer lowers its costs. When the supplier can raise IP_i above R^* in every period, under a reorder point policy IP_i is uniformly distributed on the interval $[R^* + 1, R^* + Q_r]$. In that case,

$$C_r = \frac{1}{Q_r} \sum_{c=1}^{Q_r} \bar{G}_r(R^* + c).$$

The reorder point R^* also minimizes $E[G_r(IP_i, R^*)]$, even if the supplier is unable to raise IP_i above R^* in every period: $G_p(IP_i, R^*)$ exactly compensates the retailer for any additional costs incurred due to the supplier's less than perfect reliability. Furthermore, $C_r = \min E[G_r(IP_i, R^*)]$. The function $G_p(IP_i, R^*)$ is analogous to Clark and Scarf's (1960) *induced penalty function*, because it equals the additional cost the retailer incurs when the supplier fails to deliver immediately what the retailer requests. So NC_r is a tight lower bound for the retailers' charges; from (5), our lower bound is now

$$\min E[C] = \min E[G_s(IL_s, IP, R^*)] + NC_r.$$

Turn to the supplier's charges. It is not known which policy minimizes the supplier's charges because of the location constraint: Shipping a batch to retailer i in period t might minimize expected costs in period t , but does not necessarily minimize costs in future periods. To circumvent this problem, the location constraint is relaxed, which means that batches can be freely and instantly moved from one retailer to another in each period. (This shuffling of batches occurs before demand, i.e., when the shipment release decisions are made.) Therefore, shipping a batch to a retailer with the highest need in period t does not prevent the supply chain from moving that batch to a higher need retailer in period $t + 1$.

Let \underline{C}_s be expected supplier charges with the location constraint relaxed,

$$\underline{C}_s = \min E[G_s(IL_s, IP, R^*)]$$

$$\text{s.t. } \sum_{i=1}^N IP_i \leq IL_s,$$

$$IP_i = IP_i^b + \beta_i Q_r \quad \beta_i \in \{ \dots, -1, 0, 1, \dots \}.$$

Each period the retailers' inventory positions are chosen subject to two constraints: The sum of the retailers' inventory positions cannot exceed the system's inventory level, IL_s ; and inventory can be added or subtracted from a retailer only in integer batch quantities. The first constraint applies to any feasible policy. In the second constraint, the free shuffling of inventory allows $\beta_i < 0$, whereas any feasible policy is restricted to $\beta_i \geq 0$. (Note that allowing a free shuffling of inventory among retailers is equivalent to allowing negative orders.) Thus, $\underline{C}_s + NC_r$ is a lower bound for supply chain costs, i.e.,

$$\min E[C] \geq \underline{C}_s + NC_r.$$

As with any feasible policy, a policy that yields \underline{C}_s must make two decisions: how inventory should be allocated among the retailers (what β_i to choose for each retailer), and how much inventory the supplier should order each period.

Begin with the allocation decision. Define $H(y, r)$ as the change in penalty charges when an additional batch is allocated to a retailer,

$$H(y, r) = G_p(y, r) - G_p(y + Q_r, r).$$

Note that $H(y, R^*)$ is nonincreasing in y . Therefore, penalty charges decline if $IP_j > IP_i + Q_r$ and one batch is removed from retailer j and allocated to retailer i . So the first goal of the allocation decision is to balance the retailers' inventory positions: Move a batch from the retailer with the highest inventory position to the retailer with the lowest inventory position as long as the difference between their inventory positions (before moving the batch) is greater than Q_r . After balancing the retailers' inventory positions the supplier allocates its inventory (whenever $I_s^b > 0$, otherwise the supplier has no inventory to allocate). The retailers' inventory positions are sorted in ascending order and batches are allocated in that sequence. Because $G_p(y > R^*, R^*) = 0$, the supplier stops allocating inventory when either it runs out of

inventory or when the lowest retailer inventory position is above R^* .

The supplier's ordering decision is more challenging than the allocation decision. We develop a myopic policy and then argue that the myopic policy is indeed optimal.

Define the supplier's echelon inventory position, $EI_s = IL_s + Q_r IT_s$. Recall that IL_s is measured in units and IT_s is measured in batches, so $Q_r IT_s$ is the number of units in transit to the supplier. Also define $\hat{G}_s(y, IP, R^*)$,

$$\begin{aligned} \hat{G}_s(y, IP, R^*) &= \min h_s(y - ND_r^{L_s}) \\ &+ E \left[\sum_{i=1}^N G_p(IP_i - D_r^{L_s} + \beta_i Q_r, R^*) \right] \\ \text{s.t. } \sum_{i=1}^N IP_i - D_r^{L_s} + \beta_i Q_r &\leq y - ND_r^{L_s} \\ \beta_i &\in \{ \dots, -1, 0, 1, \dots \}. \end{aligned} \quad (6)$$

In words, $\hat{G}_s(EI_s, IP, R^*)$ is the supplier's expected charges in period $t + L_s$ given that batches will be allocated among the retailers to balance their inventories as much as possible. (Note that $EI_s - ND_r^{L_s}$ is the supplier's echelon inventory level in period $t + L_s$.) The supplier's period t order is $EI_s - EI_s^b$, where $EI_s - EI_s^b$ must be a nonnegative integer multiple of $Q_r Q_s$. From (6), each batch the supplier orders will certainly increase period $t + L_s$ costs by $h_s Q_r$. However, each batch the supplier orders will also reduce penalty charges. The supplier must order a multiple of Q_s batches. Consider the θ th set of Q_s batches the supplier could order, $\theta \in \{1, 2, \dots\}$. Assuming the supplier has ordered $(\theta - 1)Q_s$ batches, the supplier should order an additional Q_s batches if the expected reduction in penalty charges due to those batches is larger than the expected increase in holding costs due to those batches, $h_s Q_r Q_s$. Hence, we need to evaluate the expected reduction in penalty charges.

In period $t + L_s$ the supplier will be able to allocate to the retailers the batches it orders in period t plus the $I_s + IT_s^b$ unallocated batches that are in the system in period t . Consider the b th batch among those batches. In period $t + L_s$ it might be allocated to some retailer. If it is allocated to a retailer, then it certainly will be

allocated to the retailer with the lowest inventory position. Let $Z_b(t + L_s)$ be that retailer's period $t + L_s$ inventory position before this batch is allocated, that is, if that batch is allocated to that retailer, then that retailer's inventory position is increased from $Z_b(t + L_s)$ to $Z_b(t + L_s) + Q_r$. Thus, that batch will lower period $t + L_s$ penalty charges by $H(Z_b(t + L_s), R^*)$. Note that $H(y > R^*, R^*) = 0$ and $H(y \leq -Q_r, R^*) = (h_s + p_r)Q_r$. Hence, to minimize period $t + L_s$ charges, the supplier should order the θ th set of Q_s batches if

$$\begin{aligned} h_s Q_r Q_s &\leq (h_s + p_r) Q_r \sum_{b=I_s+IT_s^b+(\theta-1)Q_s+1}^{I_s+IT_s^b+\theta Q_s} \\ &\times \Pr(Z_b(t + L_s) \leq -Q_r) \\ &+ \sum_{w=-Q_r+1}^{R^*} \Pr(Z_b(t + L_s) = w) H(w, R^*), \end{aligned} \quad (7)$$

where the right side above is the expected reduction in period $t + L_s$ penalty charges due to ordering those batches. Because $H(y, R^*)$ is decreasing in y and Z_{j+1} stochastically dominates Z_j , there exists a $\theta^* \geq 0$ such that (7) holds for all $\theta \leq \theta^*$, but not for $\theta > \theta^*$. Thus, given that the supplier has decided to order $(\theta - 1)Q_s$ batches, it will reduce period $t + L_s$ charges if it orders one more set of Q_s batches as long as $\theta \leq \theta^*$. In short, the myopic policy orders $\theta^* Q_s$ batches.

The implementation of (7) requires the distribution function of $Z_b(t + L_s)$. Let $Y_i(IP_i, R)$ equal the minimum number of batches that must be allocated to retailer i over periods $[t + 1, t + L_s]$ so that $IP_i(t + L_s) > R$,

$$\Pr(Y_i(IP_i, R) \leq x) = \Pr(D_r^{L_s} \leq IP_i - R + x Q_r - 1).$$

Let $Y(IP, R)$ equal the minimum number of batches that must be allocated to all retailers over periods $[t + 1, t + L_s]$ so that they all have an inventory position greater than R ,

$$Y(IP, R) = \sum_{i=1}^N Y_i(IP_i, R).$$

(Demand independence across retailers implies the above is a simple convolution.) If fewer than b batches

are required to raise the inventory position of all of the retailers above R , then the b th batch is surely allocated to a retailer with an inventory position above R :

$$\Pr(Z_b(t + L_s) \geq R + 1) = \Pr(Y(IP, R) \leq b - 1).$$

We now argue that the myopic policy is optimal, i.e., yields mean cost \underline{C}_s . Since $Z_b(t + L_s)$ stochastically dominates $Z_b(t + L_s + 1)$, the θ th set of Q_s batches that the supplier could order in period t will reduce period $t + L_s + 1$ penalty charges more than period $t + L_s$ penalty charges. So, if ordering the θ th set of Q_s batches is justified in period t (i.e., period $t + L_s$ penalty charges are reduced more than the additional holding costs those batches create), then ordering them is certainly justified in any later period. In other words, the actions taken in period t to minimize costs in period $t + L_s$ do not prevent the minimization of costs in subsequent periods. Because the myopic policy in period t does not interfere with the myopic policy in future periods, the myopic policy minimizes costs in every period, i.e., it is optimal.

Although the described policy yields $E[\underline{C}_s]$ over an infinite horizon, there unfortunately does not appear to be a simple method to evaluate $E[\underline{C}_s]$ exactly. (The state space for IP is quite large, and it is not known how to evaluate the steady state distribution of IP under the myopic policy.) Therefore, as a practical solution we use simulation to estimate $E[\underline{C}_s]$. We initialize the simulation with $IP_0^b = 0$ and IP_i^b equal to a random draw from a uniform random variable over the interval $[-Q_r, 0]$, i.e., the system begins with no inventory. We begin collecting cost data L_s periods after the first supplier order and stop collecting cost data $L_s - 1$ periods after the γ th supplier order, for some $\gamma \geq 1$. Let \underline{c}_s^j be the mean cost in the j th simulation. After n simulations, let \underline{C}_s^* be the estimate of $E[\underline{C}_s]$,

$$\underline{C}_s^* = \frac{1}{n} \sum_{j=1}^n \underline{c}_s^j.$$

Assuming \underline{c}_s^j are drawn from identical and independent normally distributed random variables, \underline{C}_s^* has a

Student t distribution with $n - 1$ degrees of freedom.⁴ Hence, confidence intervals can be constructed for \underline{C}_s^* . Our estimate of the lower bound equals $\underline{C}_s^* + NC_r$.

Chen and Zheng (1994) develop several lower bounds for multi-echelon inventory systems. Their model is somewhat different: There are no order quantity restrictions and they include explicit ordering costs. Nevertheless, they also divide the system's costs into two components and evaluate lower bounds for each one. However, since their model does not include minimum order quantities, they are able to explicitly evaluate each component's lower bound.

6. Full Information Policy

Full information provides the supplier with data to: (1) improve its order quantity decisions and (2) to improve its allocation decisions. It can make better orders because the supplier's local inventory data are not a perfect proxy for the supply chain's replenishment need: They do not measure the supply chain's total inventory nor do they capture how the supply chain's inventory is allocated among the retailers. Allocation of inventory can be enhanced for two reasons: (1) allocations can be based on the retailers' inventory positions rather than the number of batches they order (when Q_r is large each retailer might order one batch, but their inventory positions may differ substantially); and (2) the allocation of a batch can be based on the retailers' inventory positions in the period the batch is shipped rather than in the period the batch is ordered (as is done with batch priority).

Although the lower bound decision rule (7) is not optimal if batches cannot be freely moved among retailers, it does do, at least in part, what a full information policy should do: Its order recommendation is influenced by both the total amount of inventory in the system as well as how that inventory is distributed within the system. Since that rule is reasonable, we assume that under full information the

⁴ For each scenario we tested in the numerical study we evaluated the kurtosis of the data. (The kurtosis of a normal random variable is zero). We constructed a QQ plot for the 10 scenarios with the highest absolute kurtosis values, which revealed that even for those scenarios it is reasonable to assume the normal distribution.

supplier uses (7) to decide its order quantity each period.

We conjecture that (7) will tend to order fewer batches than optimal because it underestimates the value of inventory for reducing penalty charges: The expected inventory position of the retailer allocated the b th batch in the supplier's order will be lower without the free shuffling of inventory than with the free shuffling of inventory. However, this bias occurs only if the retailers' inventory positions are out of balance, i.e., only if a free shuffling is useful. We conjecture that out of balance situations are less likely as Q_r is increased, consumer demand variability is decreased, or as L_s is decreased.

The full information policy must also specify how inventory is allocated/shipped to the retailers. The lower bound policy uses a retailer reorder point, R^* , to decide when to ship inventory to the retailers, but it is easy to confirm with a few numerical examples that R^* is generally a poor choice in actual implementation.⁵ Unfortunately, there does not appear to be a simple technique for choosing the best reorder point. Therefore, we implement the retailer's optimal reorder point under traditional information; that reorder point is substituted for R^* in (7). Note that this allocation policy is an improvement over *batch allocation* because the supplier will always ship a batch to the retailer with the lowest inventory position in the period the batch is shipped rather than to a retailer that had the lowest inventory position in some prior period (when the batch was ordered).

As with the lower bound, a priori evaluation of expected per period costs with this full information policy is not possible. Hence, we evaluate this policy using the same simulation approach used to evaluate the lower bound. There is no guarantee that this full

⁵ Intuition also suggests that R^* is not a good choice. When h_r is small relative to h_s , R^* is quite large. That means that once inventory arrives at the supplier it will be immediately shipped to the retailers. (If the supplier held the inventory, then it would fill the retailers' orders quickly and their average inventory would be quite high.) Pushing inventory down to the retailers immediately is not risky when inventory can be freely moved among the retailers, but it can be quite costly if rebalancing is not possible. At the other extreme, if h_r is large relative to h_s , then R^* will be low and the retailers will probably run higher back orders than they should.

information policy will perform better than the traditional information policy. However, the numerical study finds that it did perform better in almost all of the tested scenarios.

7. Numerical Study

This section reports on a numerical study that evaluates supply chain costs with the traditional information policy, the full information policy, and the lower bound. Section 7.1 details the parameters and methods and §7.2 details the results.

7.1. Parameters and Methods

The numerical study consists of the 768 scenarios formed from all combinations of the following parameters:

$$N \in \{4, 16\} \quad Q_r \in \{2, 4, 8, 16\} \quad L_r \in \{1, 5\}$$

$$h_r = 1 - h_s \quad p_r \in \{5, 25\} \quad \sigma_r \in \{0.36, 1\}$$

$$Q_s \in \{1, 4, 16\} \quad L_s \in \{1, 5\} \quad h_s \in \{0.5, 1\}.$$

In all scenarios $\mu_r = 1$. The parameters are chosen to reflect a wide range of situations: small and large batch sizes (relative to mean demand), short and long lead times, cheap and expensive supplier inventory, low and high back-order penalty costs, and low and high consumer demand variability. When $\sigma_r = 1$, demand at each retailer is Poisson. The distribution is truncated, so that $\Pr(D_r^1 \leq 7) = 1$, whereas in the actual distribution $\Pr(D_r^1 \leq 7) = 0.99999$.⁶ When $\sigma_r = 0.36$, demand is a discretization of a Normal distribution with mean 1 and standard deviation $\frac{1}{3}$:

$$\Pr(D_r^1 = 0) = 0.066807; \quad \Pr(D_r^1 = 1) = 0.866386;$$


$$\Pr(D_r^1 = 2) = 0.066807; \quad \Pr(D_r^1 \geq 3) = 0.$$

Although we do not expect to observe the above distribution in practice, it provides a representation of supply chain behavior with low consumer demand variability.

For each scenario the system-wide cost-minimizing reorder points are found assuming traditional information. The costs evaluations are exact when R_s

⁶ The mean of that distribution is 0.99999 and the standard deviation is 0.99992.

Figure 1 Traditional vs. Full Information



≥ -1 , and approximate otherwise. We confirmed via simulation that the approximation introduces little error.

To estimate the full information policy cost and the lower bound we conducted 40 simulations of each scenario, where each scenario collected data over 10 supplier order cycles ($\gamma = 10$). Due to sampling error, it is possible to observe a mean lower bound that is greater than the mean full information policy cost. For those scenarios in which that occurred, we disposed of the original data and conducted 90 additional simulations per scenario. These latter data are used in our analysis.⁷ For each scenario we constructed 95% confidence intervals for both the lower bound and the full information policy cost. The mean confidence interval, expressed as a percentage of the mean estimate, is 0.5% for the lower bound and 0.8% for the full information policy cost. We conclude that the estimates are sufficiently accurate.

7.2. Results

Figure 1 displays a histogram of the cost difference between the traditional information policy and the full

information policy, i.e., our realized value of information sharing. Those differences, as well as the cost differences displayed in the other figures, are all expressed as percentages of the traditional information policy cost. The mean benefit of the full information policy is 2.2%, the maximum benefit is 12.1%, and for 95% of scenarios the percentage is no more than 5.6%.

Figure 2 displays a histogram of the cost difference between the full information policy and the lower bound, i.e., the potential cost improvement if the optimal full information policy were implemented. The mean is 1.2%, the maximum is 10.5%, and for 95% of scenarios the percentage is no more than 4.1%. (In a few scenarios the lower bound estimate exceeds the feasible policy estimate, which we attribute to sampling error.) We conclude that the performance of our full information policy is quite close to optimal.

Figure 3 presents the upper bound on the value of shared information: the cost difference between the traditional information policy and the lower bound. The mean value is 3.4%, the maximum is 13.8%, and for 95% of scenarios the value is no more than 8.0%.

We investigated two other possible benefits of information technology. One is that information technology reduces order processing times, which we model as a reduction in L_r . The second benefit is that information

⁷ Due to this procedure there is a positive bias for the difference between the full information cost and the lower bound cost in the scenarios that are not discarded.

Figure 2 Full Information vs Lower Bound





Figure 3 Maximum Value of Shared Information



technology reduces order processing costs. We model that as a reduction in Q_r , because a firm can justify more frequent ordering with lower order processing costs. Although the magnitude of these effects will depend on the context, we consider the impact of cutting batch sizes in half or the impact of reductions in L_r from 5 to 3, from 3 to 1, and from 5 to 1. (As we mention in the introduc-

tion, we feel that these are reasonable changes to the L_r and Q_r parameters.) In each scenario we assume traditional information sharing. Table 1 displays data on the potential cost reductions. Across all scenarios it is apparent that a reduction in L_r or Q_r can have a significant impact on supply chain costs and a significantly greater impact than sharing information.

Table 1 Shared Information vs. Lead Time Reduction or Batch Size Reduction

Upper Bound on the Value of Complete Information*		% Decrease in Total Supply Chain Cost						
		Retailer Lead Time Reduction With Traditional Information			Retailer Batch Size Reduction With Traditional Information			
		3 to 1	5 to 3	5 to 1	16 to 8	8 to 4	4 to 2	2 to 1
Minimum	0%	5%	4%	9%	16%	8%	3%	1%
Mean	3%	29%	21%	43%	36%	27%	17%	10%
Median	3%	30%	22%	46%	36%	26%	15%	7%
Maximum	14%	57%	35%	71%	49%	47%	43%	37%

* (Traditional Information Cost – Lower Bound)/Traditional Information Cost.

8. Discussion

Our results are surprising. Indeed, we undertook this research with the strong expectation that we would be able to demonstrate significant benefits to information sharing in these models. So why do the data speak the opposite conclusion? We conjecture the following answer: The retailers' orders convey to the supplier a substantial portion of the information the supplier needs to perform its ordering and allocation functions. When a retailer is flush with inventory, its demand information provides little value to the supplier because the retailer has no short-term need for an additional batch. A retailer's demand information is most valuable when the retailer's inventory approaches a level that should trigger the supplier to order additional inventory, but this is also precisely when the retailer is likely to submit an order. Hence, just as the retailer's demand information becomes most valuable to the supplier, the retailer is likely to submit an order, thereby conveying the necessary information without explicitly sharing demand data.

How can we reconcile our results with the clear trend across many industries to apply information technology to logistics and inventory management? In fact, our results are quite consistent with that trend: We do find substantial savings from lead time and batch size reductions, both of which are facilitated by the implementation of information technology. We only conclude that the observed benefits of information technology in practice are due more to the impact

of information technology on lead time and batch size than in facilitating information sharing. Further, because we measure the value of information sharing using a lower bound on any feasible inventory policy that uses shared information, this finding will not be changed by the formulation of more sophisticated algorithms for sharing information.

Although our model is representative of many actual supply chains, we recognize that our conclusion is limited to the setting we consider. In particular, in our model demand is known, the retailers are identical, there exists only a single source for inventory, there are no capacity constraints, there are no incentive conflicts among the supply chain's firms, and firms choose rational ordering policies.

We anticipate that information sharing can have a significantly greater value in environments with unknown demand, for example, early sales of new products or established products on promotion. In those settings information sharing would improve the supplier's ability to detect shifts in the demand process.

We study a model with identical retailers for two reasons. First, nonidentical retailer models are far more complex to analyze. With nonidentical retailers it is not even known how to evaluate reorder point policies with traditional information. Second, in our setting we anticipate that information has the highest value with identical retailers. Information sharing allows the supplier to identify which retailers have the highest need for replenishments. This is

most important when the retailers are nondistinguishable, i.e., when they have identical demand processes. In fact, Aviv and Federgruen (1998) obtain comparable results to ours in a model with nonidentical retailers. Our justifications notwithstanding, additional research is needed to assess information sharing with nonidentical retailers.

Although a single inventory source is reasonable for many settings, in some supply chains the firms may have access to a second inventory source, albeit at a higher cost. In those settings information sharing may allow the supply chain to better decide when it should utilize its alternative sources.

Gavirneni et al. (1999) found that information sharing is most valuable when capacity is not restrictive; information is valuable only if the system has the flexibility to respond to the information. Hence, imposing a capacity constraint on the supplier would probably lower the value of information in our model.

We have assumed that a benevolent dictator decides all inventory shipments. This is reasonable when the sole objective is minimizing total supply chain cost, in other words, it doesn't matter which firm makes the decision. However, in actual supply chains the firms might not share the same objectives. Additional research is needed to determine how the firms will behave in those settings. For example, they may needlessly hoard inventory, thereby raising costs for everyone. It is important to determine which conditions provide the players with the incentive to truthfully reveal their private information and whether revealing information eliminates destructive gaming.⁸

⁸ For many helpful comments, the authors thank Lee Schwarz, Sridhar Tayur, Paul Zipkin, and seminar participants at the Department of Industrial Engineering and Engineering Management at Stanford University and the Operations Research Center at the Massachusetts Institute of Technology. The assistance of the associate editor and the reviewers is graciously acknowledged.

References

Anand, K., H. Mendelson. 1997. Information and organization for horizontal multimarket coordination. *Management Sci.* **43** 1609–1627.

- Aviv, Y. 1998. The effect of forecasting capability on supply chain coordination. Working paper, Washington University, St. Louis, MO.
- , A. Federgruen. 1998. The operational benefits of information sharing and vendor managed inventory (VMI) programs. Working paper, Washington University, St. Louis, MO.
- Axsäter, S. 1993. Exact and approximate evaluation of batch-ordering policies for two-level inventory systems. *Oper. Res.* **41** 777–785.
- . 1997. Exact analysis of continuous review (R, Q)-policies in two-echelon inventory systems with compound Poisson demand. Working paper, Lund University, Lund, Sweden.
- Bourland, K., S. Powell, D. Pyke. 1996. Exploiting timely demand information to reduce inventories. *European J. Oper. Res.* **92** 239–253.
- Cachon, G. 1995. Exact evaluation of batch-ordering policies in two-echelon supply chains with periodic review. *Oper. Res.* Forthcoming.
- , M. Fisher. 1997. Campbell Soup's continuous product replenishment program: Evaluation and enhanced decision rules. *Prod. Oper. Management* **6** 266–276.
- , M. Lariviere. 1997. Contracting to assure supply: How to share forecasts in a supply chain. Working paper, Kellogg Graduate School of Management Northwestern University, Evanston, IL.
- Chen, F. 1997. Optimal policies for multi-echelon inventory problems with batch ordering. Working paper, Columbia University, New York.
- . 1998. Echelon reorder points, installation reorder points, and the value of centralized demand information. *Management Sci.* **44** S221–S234.
- , R. Samroengraja. 2000. A staggered ordering policy for one-warehouse multi-retailer systems. *Oper. Res.* **48**(2) 281–293.
- , Y. S. Zheng. 1994. Lower bounds for multi-echelon stochastic inventory systems. *Management Sci.* **40** 1426–1443.
- , —. 1997. One-Warehouse multi-retailer systems with centralized stock information. *Oper. Res.* **45** 275–287.
- Clark, A. J., H. E. Scarf. 1960. Optimal policies for a multi-echelon inventory problem. *Management Sci.* **6** 475–490.
- Clark, T., J. Hammond. 1997. Reengineering channel reordering processes to improve total supply chain performance. *Prod. Oper. Management* **6** 248–265.
- Graves, S. 1996. A multiechelon inventory model with fixed replenishment intervals. *Management Sci.* **42** 1–18.
- . 1999. A single-item inventory model for a nonstationary demand process. *Manufacturing Ser. Oper. Management* **1** 50–61.
- Gavirneni, S., R. Kapuscinski, S. Tayur. 1999. Value of information in capacitated supply chains. *Management Sci.* **45** 16–24.
- Kurt Salmon Associates Inc. 1993. Efficient Consumer Response: Enhancing Consumer Value in the Grocery Industry. Food Marketing Institute, Washington, DC.
- Lee, H., K. Moynadeh. 1986. Batch size and stocking levels in

- multi-echelon repairable systems. *Management Sci.* **32** 1567–1581.
- , P. Padmanabhan, S. Whang. 1997a. Information distortion in a supply chain: The bullwhip effect. *Management Sci.* **43** 546–558.
- , K. So, C. Tang. 2000. The value of information sharing in a two-level supply chain. *Management Sci.* **46**(5) 626–643.
- Liljenberg, Peter. 1996. The value of centralized information in a two-echelon inventory system with stochastic demand. Working paper, Lund University, Lund, Sweden.
- McGavin, E., L. Schwarz, J. Ward. 1993. Two-Interval inventory allocation policies in a one-warehouse N -identical retailer distribution system. *Management Sci.* **39** 1092–1107.
- Svoronos, A., P. Zipkin. 1988. Estimating the performance of multi-level inventory systems. *Oper. Res.* **36** 57–72.

Accepted by Luke Van Wassenhove; received March 13, 1997. This paper was with the authors 12 months for 2 revisions.