Optimizing Delivery Lead Time / Inventory Placement in a Two-Stage Production/Distribution System[1]

Dawn Barnes–Schuster[*]

Rötelsteig 9, CH-8037 Zürich

Yehuda Bassok[†]

Marshall School of Business

University of Southern California, Los Angeles, CA 90089

Ravi Anupindi[‡]

Stern School of Business

New York University, New York, NY 10012

First version: May 2000

Last revised: March 22, 2002

---

[1] An early version of this paper was titled “Negotiating Delivery Lead Time in a Two-Stage Decentralized System”, October 1995.

[*] Email: `dawn@barnes-schuster.ch`; Phone: 41-1-364-5085

[†] Email: `bassok@bus.usc.edu`; Phone: (213)821-1140

[‡] Email: `anupindi@stern.nyu.edu`; Phone: (212)998-0487

Abstract

In this paper we study a system composed of a supplier and buyer(s). We assume that the buyer faces random demand with a known distribution function. The supplier faces a known production lead time. The main objective of this study is to determine the optimal delivery lead time and the resulting location of the system inventory. In a system with a single supplier and a single buyer it is shown that system inventory should not be split between a buyer and supplier. Based on system parameters of shortage and holding costs, production lead times, and standard deviations of demand distributions, conditions indicating when the supplier or buyer(s) should keep the system inventory are derived. The impact of changes to these parameters on the location of system inventory is examined. For the case with multiple buyers, it is found that the supplier holds inventory for the buyers with the *smallest* standard deviations, while the buyers with the *largest* standard deviations hold their own inventory.

Keywords: Supply Chain Management, Delivery Lead Time, Inventory, Production

1 Introduction

In 1994, Compaq was able to cut its supply costs by $600 million, in part, by revamping its supply chain. Changes included clearing most of the parts inventories off the factory floor in Houston, Texas and into a nearby warehouse leased by 35 major Compaq suppliers, all of whom moved their inventories to Houston at Compaq's behest. Compaq also convinced seven sheet-metal suppliers to re-locate to Texas, and subsequently to truck the parts to Compaq exactly when and where they were needed (Henkoff, 1994).

Essentially, Compaq reduced the delivery lead times offered to it by its suppliers to zero bringing substantial savings for itself. Moreover, Compaq's suppliers were likely compensated (monetary or otherwise) as is evidenced by their willingness to relocate. Hence, the supply chain as a whole seems to have benefited from the changes. It is becoming increasingly common that suppliers relocate closer to their buyers and thereby reduce delivery lead time to zero. When is it optimal for the system to reduce delivery lead time to zero? In this paper we develop a model and present an analysis to answer this question.

Specifically, we analyze a model of a two-stage supply chain consisting of a single supplier and multiple buyers. We model *delivery lead time* as the time taken by an upstream stage to fulfill an order from the downstream stage. We let the delivery lead time be a decision variable and evaluate its effect on system-wide costs. Not surprisingly, delivery lead time and the safety inventory in the system are closely linked. We show that it is optimal for either the supplier or the buyer, but not both, to carry the safety inventory for the system. Anecdotal evidence on implementation of Just-in-Time systems suggest that buyers push safety inventory to the suppliers and expect deliveries as and when needed. We show that a policy that pushes the entire system inventory to either the supplier or the buyer, is optimal.

We first consider a two-stage serial system with a single-supplier and a single-buyer. Subsequently, we extend the analysis to a two-stage divergent system with a single-supplier and multiple-buyers. Simpson (1958) was the first to consider the question of inventory placement in a serial system. He examined serial systems where the consecutive stages could, potentially, be separated by inventory. The level of inventory between stages depended on the delivery lead time[2] quoted by the upstream stage to the downstream stage, and a given desired service level. When the service level is 100% (i.e., that an upstream stage *must* supply a downstream stage within the quoted quoted delivery lead time), he showed that the cost minimizing optimal delivery lead time an upstream stage quotes would be either zero or its total production and transportation time. Our single-buyer single-supplier system is closely related to a two-stage version of Simpson's serial system. In Simpson's model inventory between consecutive stages could be carried only at one of the stages. In addition, Simpson uses a service level approach to meeting the delivery lead time and minimizes holding, order, and variable costs. In contrast, we allow inventory to be potentially carried at both stages. Instead of a service level, we allow the supplier to incur expediting costs to guarantee a delivery lead time.[3] Inderfurth (1994) extended Simpson's work to divergent inventory systems. That is, he considered systems where parts can go into more than one type of finished product as they move downstream. He showed, once again, that the delivery lead time should be set to either zero or the production/transportation time. He, however, assumed that the delivery lead times

---

[2] They call it the service time.

[3] Simpson's model can be modified to mimic our situation. Since we allow for inventory to be held either at the buyer or the supplier, we need to add a dummy transportation stage between the buyer and the supplier in Simpson's model. In addition, we need to add a shortage cost to Simpson's average cost function.

quoted to all downstream stages are *identical*. Our model differs from Inderfurth's in ways similar to those identified between our model and Simpson's. In addition, we show that identical delivery lead times need not be optimal for divergent systems. Graves and Willems (1998) extend Inderfurth's work to a more general model of spanning trees. Their intent is to create a tactical tool for industry in its efforts to model and improve supply chains. Unlike Inderfurth, they allow for non-identical delivery lead times for various downstream stages. Like Inderfurth, they ignore the possibility of shortages, though potentially in a different way. They optimize the delivery lead times based on demand equaling its upper bound.

In addition, there is a large body of literature that covers multi-stage systems starting with the pioneering work of Clark & Scarf (1960); also see Eppen & Schrage (1981), Federgruen & Zipkin (1984), Rosling (1989), Langenhoff & Zijm (1990), Glasserman & Tayur (1995), and Gallego & Zipkin (1999). All of these papers study a multi-stage system from the perspective of a single decision maker. Decentralized approaches to the study of multi-stage systems include Lee & Whang (1999) and Cachon & Zipkin (1999). Furthermore, the delivery lead time (or its distribution) is an exogenous variable in all of these studies. In contrast, we endogenize the delivery lead time variable and find its optimal value. This deviation is not inconsequential. To see this point clearly, we compare our two-stage serial model with the celebrated model of Clark and Scarf. In Clark and Scarf the delivery lead time is the transportation time. Consequently, the time products spend in transport cannot be reduced by inventory. In contrast, we assume that the delivery lead time includes the production lead time of the upstream stage. This implies that the supplier can reduce the delivery lead time by keeping an appropriate level of inventory. In Clark and Scarf the supplier delivers only based on its on-hand inventory. This implies that an order may delivered to the buyer over multiple shipments. In contrast, we allow for the supplier to expedite production[4], if necessary. Hence she is always able to satisfy the buyer's entire demand. To this extent, our model is similar in spirit to those considered by Inderfurth and Graves and Willems.

We examine the following two-stage models: a single-supplier single-buyer system and a single-supplier multiple-buyers system. We assume that a buyer follows a periodic review policy. Using this policy, at the beginning of every period, the buyer observes his own inventory and then places an order. The order is delivered (in its entirety) by the supplier after a

---

[4] This can also be achieved through subcontracting remaining work on the work-in-process inventory.

known, but negotiated, delivery lead time. It is well known (Arrow, Karlin & Scarf, 1958) that it is optimal for the buyer to raise his inventory position to a base-stock level.

We assume that the supplier also follows a periodic review policy and knows the distribution of the orders placed by the buyer. Observe that the distribution of the order process to the supplier is identical to the distribution of the demand process faced by the buyer when the buyer follows a base stock policy. We assume that the supplier's *production lead time* is known and independent of order size. This is a standard assumption made by a majority of the work in multi-stage systems. Finally, we assume that the supplier will always be able to satisfy the buyer's order within the delivery lead time with expediting, if necessary. We assume that the marginal expediting cost is proportional to the remaining work content; the exact details of this cost structure are discussed later in Section 2.2. Our supplier's model is justified when she has sufficient capacity to meet orders by expediting, if necessary. A supplier with a high degree of flexibility will also satisfy the assumption. Several researchers have made the infinite capacity assumption where the supplier is able to deliver the retailer's entire order in a short time period at a fixed marginal cost; see for example Lee et. al (1997) and Chen et. al. (2000). In contrast, our expediting cost model is more reasonable. Inderfurth (1994) and Graves and Willems (1998) also make the same assumption but do not allow for expediting. Consequently, the supplier in their model has to keep inventory to respond to the upper support of the demand distribution. In contrast, we allow the supplier to expedite production, as necessary to meet the buyer's demand in entirety. Thus while the supplier in Inderfurth (1994) and Graves and Willems (1998) responds to orders by a pure safety-stock strategy, in our model, she uses a mix of safety stock and safety capacity to respond to orders within a delivery lead time. It is straightforward to see that their model of the supplier is a special case of ours.

How does the delivery lead time affect the operations of this two-stage model? In a pure make-to-stock system, the supplier carries all the inventory and hence experiences all the demand risk. She can reduce the impact of this demand uncertainty increasing the delivery lead time to the buyer. To see this point, suppose the buyer places an order at time t, and this order will be delivered at time $t + l_d$, where l_d is the delivery lead time. Suppose that $l_d \leq l_p + 1$, where l_p is the production lead time.[5] At time $t + l_d$ the order is withdrawn from

---

[5] We assume that the delivery lead time is bounded above by the lead time of a pure make-to-order system. A

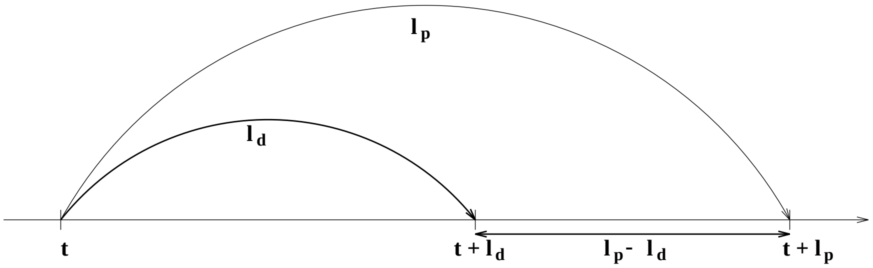

Figure 1: Effect of Delivery Lead Time on Supplier's Risk

the supplier's inventory and delivered to the buyer. At time $t + l_p$ production of the order is complete. Note that the safety stock that the supplier keeps is then proportional to $l_p - l_d$ (an exact expression for this quantity will be given later). By examining Figure 1, it is easy to see that decreasing the delivery lead time, l_d, increases $l_p - l_d$, resulting in a longer period of time for which the supplier must hold safety stock. Hence some of the risk has been transferred from the buyer to the supplier. Certainly, the opposite is true should the delivery lead time be increased.[6]

We are interested in determining the cost-effective delivery lead time for a two-stage system consisting of a supplier and buyer(s). We assume that the players in the supply chain operate as a team; that is, they minimize the total expected cost incurred by the system. Such an objective is justified when the system represents one firm or when the buyer and the supplier have an agreement to share any savings resulting from such a cooperative arrangement. The main contributions of this paper are as follows.

1. In general, we show that it is never optimal for the safety stock of the system to be split between an individual buyer and the supplier, but should be carried by one or the other.
2. For a single-supplier and a single-buyer system, we develop conditions that indicate who should carry the system inventory and hence the delivery lead time offered.
3. For a single-supplier and two-buyer system, we show that either *both* buyers keep their own safety stock or the supplier keeps the safety stock for both buyers.

---

larger delivery lead time gives the supplier no added advantage.

[6] Clearly decreasing the production lead time should be a primary mechanism for lowering costs. We assume that operational improvements have already been made to minimize the production lead time and focus on lead time coordination mechanisms to improve the supply chain performance.

4. For a single-supplier and multiple-buyers (more than two) system, we develop conditions for when the supplier will hold inventory for all the buyers, when she will hold inventory for only a subset of the buyers and when all the buyers will hold their own inventory. These conditions depend on both the cost and the demand variability parameters. When the supplier holds the system safety stock for only a subset of the buyers, surprisingly it is for those with the smallest standard deviations of demand.

5. We give comparative statics results on how the location of inventory changes with the cost parameters.

6. Furthermore, we show an important stability result in which if the current contractual arrangements require that the supplier carry inventory for some of the buyers, then contracting with new buyers will not result in an unfavorable change in these arrangements with the existing buyers. That is, the supplier will continue to carry inventory for those buyers for whom she currently does and may, in addition, start carrying inventory for those she earlier did not.

The rest of the paper will proceed as follows. In Section 2, we discuss the single supplier-single buyer model and results. We extend this model in Section 3 to a single supplier-multiple buyer model. Finally we conclude in Section 4. All proofs are given in the Appendix.

2 A Single-Supplier Single-Buyer Model

Consider a periodic review single-supplier single-buyer system. The buyer faces random demand in each period, independent and identically distributed with probability density and distribution functions, respectively, $f(\cdot)$ and $F(\cdot)$. Both the supplier and the buyer follow a stationary base-stock policy with base stock levels of S^s and S^b, respectively. Both incur a unit holding cost for leftover goods at the end of the period, denoted by h^s and h^b respectively. The buyer incurs a unit penalty cost, p^b, for excess demand not satisfied. The supplier's normal production lead time is l_p and she quotes a delivery lead time of l_d to the buyer. That is, the buyer's order in period t is delivered in period $t + l_d$. The supplier always satisfies the demand of the buyer. If the buyer's order is less than the supplier's base stock level, S^s, it is satisfied completely from stock. Otherwise, the order quantity in excess of S^s is satisfied by expediting her production pipeline stock. We assume that the cost of expediting one unit by one period is p^s. Thus, the total cost of meeting excess order by expediting is proportional to the remaining work on the work-in-process inventory. The exact mechanism of expediting for a specific order is outlined in more detail later when we develop the expected cost function of the supplier. For the rest of the paper, we will merely refer to p^s as the supplier's penalty cost. Without loss of generality, we ignore any cost of materials and normal production since all demand is ultimately satisfied. To make the problem interesting, we assume that the unit penalty/expediting cost is larger than the corresponding unit holding cost of the buyer / supplier.

The timing of events within a period, say t, is as follows:

1. Buyer observes the demand in period t and places an order equal the demand observed.
2. Supplier delivers the goods the buyer ordered in period $t - l_d$.
3. Buyer receives goods which were ordered in period $t - l_d$ and satisfies as much of the demand in period t as is possible.
4. Supplier completes production of order begun in period $t - l_p$ and initiates production on order placed by the buyer in period t.
5. On-hand and on-order inventories are counted and costs calculated.

Note that $l_d = 0$ corresponds to the situation of Vendor Managed Inventory (VMI)[7]. That is, when demand materializes, the goods are removed from the supplier's inventory, instantaneously passed through the buyer's inventory and passed on to the customer. When $l_d = 1$, we model the standard infinite horizon newsvendor model with zero lead time. That is, the buyer must satisfy demand from inventory before placing a replenishment order (which still arrives immediately).

Recall that, because of the supplier's capability to expedite production when necessary, the system's cost function is separable in the buyer's and supplier's stocking decisions, and as such, the buyer's and supplier's stocking problems can be solved independently.

---

[7] Under VMI, the supplier makes the stocking decisions for her product on the buyer's shelves.

2.1 Buyer's Model

The buyer's optimization problem is an infinite horizon standard inventory problem with a fixed lead time. We assume that the supplier is capable of supplying, in a single shipment, any quantity the buyer asks for after a fixed delivery lead time (see discussion in Section 2.2). It is well known that this problem can be reduced to a problem with zero lead time by adjusting the demand distribution to account for the lead time demand (Arrow, Karlin & Scarf, 1958), for which a base stock policy is optimal. Let us denote the solution of the infinite horizon average cost problem and the resulting stationary base stock level, S^b, by

$$F_{l_d}(S^b) = \frac{p^b}{p^b + h^b} \tag{1}$$

where $F_{l_d}(\cdot)$ is the l_d-fold convolution of F with itself for all $l_d \geq 1$. (Should $l_d = 0$, then the buyer is capable of acquiring an order with no wait, and so gives his customers a 100% service level while holding no inventory.)

Using this stationary policy, it is easy to write an explicit expression for the expected average cost during a period. Each period, the buyer will pay an average of *variable cost* × *average periodic demand* to acquire the goods. As this will remain a constant in the cost equation, we will ignore its presence from here on out. Hence, we will concern ourselves only with the expected holding and shortage costs for following the policy of ordering up to S^b.[8]

Let x represent the single period demand random variable, whose mean is μ. The buyer's average inventory cost per period, $C^b(l_d)$, can be written as

$$\begin{aligned} C^b(l_d) &= h^b \int_0^{S^b} \left(S^b - x\right) f_{l_d}(x)dx + p^b \int_{S^b}^{\infty} \left(x - S^b\right) f_{l_d}(x)dx \\ &= (h^b + p^b)S^b F_{l_d}(S^b) - p^b S^b + p^b \int_0^{\infty} x f_{l_d}(x)dx - (h^b + p^b) \int_0^{S^b} x f_{l_d}(x)dx \\ &= p^b l_d \mu - (h^b + p^b) \int_0^{S^b} x f_{l_d}(x)dx \end{aligned} \tag{2}$$

where $f_0(\cdot) \equiv 0$. When the single period demand distribution is Normal with standard deviation σ, we have

$$S^b = l_d \mu + k^b \sqrt{l_d}\, \sigma \tag{3}$$

---

[8] We assume that any purchasing cost the buyer or supplier must pay to raise their initial inventory on-hand and in process to some S^j for $j \in \{s, b\}$ is to be treated as sunk cost. It is not likely that the system would start empty, and the difference in costs from either raising the inventory to this S^j or letting it fall to S^j would only change the results quantitatively, and not qualitatively.

where $k^b = \Phi^{-1}\left(\frac{p^b}{p^b+h^b}\right)$ and $\Phi(\cdot)$ is the standard normal distribution.

2.2 Supplier's Model

Now we will develop the expected cost function for the supplier. We make the assumption here that the supplier will always satisfy all of the buyer's demand. The supplier achieves this in the following manner. Suppose that in period t, d units of goods are to be delivered to the buyer. The supplier satisfies as much of the demand as possible from her on hand inventory, I_t. If there is still excess demand to be satisfied, she expedites production of pipeline stock, at an additional cost of p^s per unit, for as many of the goods which were to have been finished in period $t+1$ as she needs and completes them in period t. If this is still not enough to cover all of the demand in period t, she expedites production for the needed items from the goods that were to have been finished in period $t+2$, at an additional cost of $2p^s$ per unit. She continues expediting pipeline stock until she can fill all of the demand, beginning from raw materials if necessary. Thus the unit expediting cost is proportional to the number of periods expedited. It is reasonable to expect that the work content in each time period is identical. Thus the unit expediting cost is proportional to the remaining work content. Of course, we assume that the supplier has sufficient capacity to expedite the production as necessary.

The supplier also follows an order-up-to policy. Since the supplier satisfies demand from stock or by expediting the pipeline, the supplier's production system can be rewritten as a system with $l_p - l_d$ as its effective production lead time. This is easily seen since the supplier, with l_p periods of production lead time, begins production on the goods which will replenish the order the period it is received, but doesn't deduct the goods from inventory and deliver them to the buyer until l_d periods later. Production of the goods to replenish the order finish l_p periods after they were begun and are themselves available to satisfy other orders the period after they finish production. Thus the supplier is keeping inventory for the last $l_p - l_d$ production periods, and not the whole l_p periods. Hence, like the buyer, the supplier produces up to S^s, where

$$F_{l_p - l_d + 1}\left(S^s\right) = \frac{p^s}{p^s + h^s}. \tag{4}$$

Now, we develop a cost function for the supplier. It is easy to see that the expediting costs in our *supplier's model* are calculated in much the same fashion as the shortage costs in the *standard inventory model* with fixed lead time and backlogging. Consider the demand in period t, d_t, and the inventory on-hand at the supplier in period t, I_t. Clearly, if there is enough inventory on hand to satisfy the demand, then no shortage (or expediting) costs will be charged, regardless of which model is being considered. So, consider when $d_t > I_t$. Then, the supplier is short by $S_t^1 = d_t - I_t$. Suppose that an amount, O_{t+1}, from the supplier's production pipeline, is due to arrive the following period. Either this will be sufficient to cover S_t^1 or it won't suffice. If $O_{t+1} > S_t^1$ then all shortages for period t will be covered and the only shortage (expediting) costs charged for demand d_t will be for S_t^1 which will be paid in period t. If $O_{t+1} < S_t^1$ then the supplier needs to consider O_{t+2} and expedite from there. This continues until sufficient goods are expedited to satisfy all the demand. Thus the only difference between the two models is that in a standard inventory model, backlog satisfaction is postponed until the period the goods become available, while in our model, expediting results in immediate satisfaction of backlogged demand. Therefore, as holding are charged identically in both systems, the only difference in the cost structure of the two models lies in the allocation of shortage costs to periods. Hence, while the actual periods of demand satisfaction differ, the total cost for not being able to satisfy all of the demand d_t in period t remains the same, regardless of which model is used.

As in the case of the buyer, we will choose to ignore the per unit variable cost of production as it will only result in a constant added to the average cost equation. Hence, the supplier's average cost per period can be written as

$$C^s(l_d) = p^s(l_p - l_d + 1)\mu - (h^s + p^s)\int_0^{S^s} x f_{l_p - l_d + 1}(x)dx \tag{5}$$

Once again, when demand is normally distributed, S^s can be written as

$$S^s = (l_p - l_d + 1)\mu + k^s\sigma\sqrt{l_p - l_d + 1}. \tag{6}$$

2.3 Optimal Delivery Lead Time

We can now develop a formulation for the expected system (buyer and supplier) cost, and hence find the optimal delivery lead time.

Lemma 1 *The expected periodic cost function for a system with one supplier and one buyer* *with normally distributed demand can be written as*

$$C(l_d) = \sigma\left((h^s + p^s)\phi(k^s)\sqrt{l_p - l_d + 1} + (h^b + p^b)\phi(k^b)\sqrt{l_d}\right) \tag{7}$$

Moreover, this function is concave in l_d with a minimum occurring at one of the extrema of the function where $l_d \in \{0, l_p + 1\}$.

The extreme $l_d = 0$ corresponds to the supplier keeping all of the inventory while the extreme $l_d = l_p + 1$ corresponds to the buyer keeping all of the inventory. Specifically, we need to compare the ratio of costs at the two extremes (cost for inventory at the buyer to cost for inventory at the supplier) which can be written as follows:

$$\frac{C(l_p+1)}{C(0)} = \frac{\sigma(h^b + p^b)\phi(k^b)\sqrt{l_p+1}}{\sigma(h^s + p^s)\phi(k^s)\sqrt{l_p+1}} = \frac{(h^b + p^b)\phi(k^b)}{(h^s + p^s)\phi(k^s)} \tag{8}$$

Exactly which of these two choices is best will depend on the values of the cost parameters and the following result follows immediately.

Proposition 1 *Let β be the ratio of expected marginal holding and penalty cost of the buyer to the expected marginal holding and penalty cost of the supplier at optimality. That is,*

$$\beta = \frac{(h^b + p^b)\phi(k^b)}{(h^s + p^s)\phi(k^s)}. \tag{9}$$

If $\beta < 1$, then the buyer keeps the system inventory. If $\beta > 1$ then the supplier keeps the system inventory. If $\beta = 1$ then the inventory could be kept at either the buyer or the supplier.

Observe that the decision is independent of σ and l_p but depends on the respective shortage / expediting and holding costs via β. The following Proposition develops the comparative statics of the optimal inventory location based on the various problem parameters.

Proposition 2 *Consider the four cost parameters in the model – the holding and penalty costs for the supplier (h^s and p^s) and for the buyer (h^b and p^b). For any given values of three of these parameters, there exists a unique threshold value of the fourth parameter that determines the location of inventory in the system as outlined below:*

Parameters		*Location of Inventory*	
Fixed	*Threshold*	*Buyer*	*Supplier*
h^s, p^s, *and* h^b	p^{b*}	$\forall\ p^b < p^{b*}$	*otherwise*
h^s, p^s, *and* p^b	h^{b*}	$\forall\ h^b < h^{b*}$	*otherwise*
h^s, h^b, *and* p^b	p^{s*}	$\forall\ p^s \geq p^{s*}$	*otherwise*
p^s, h^b, *and* p^b	h^{s*}	$\forall\ h^s \geq h^{s*}$	*otherwise*

Observations

From any set of parameters for which the buyer holds the inventory, the threshold curve can be crossed by sufficiently increasing the buyer's shortage or holding cost, or by sufficiently decreasing the supplier's shortage or holding cost. The case for the holding costs should seem intuitively clear. Sufficiently increasing the buyer's holding cost or sufficiently decreasing the supplier's holding cost should cause the system to consider switching its inventory to the cheaper locale.

On the other hand, the change resulting from increasing the buyer's shortage cost or decreasing the supplier's expediting cost is not as intuitively obvious. Let us first consider the case of increasing the buyer's shortage cost. Intuitively, one might think that the buyer should keep the system inventory since the consequences of a shortage are large. But our results indicate the opposite. Notice that the optimal (system) inventory that the buyer keeps, S^b, and the total cost at optimality, is increasing in the buyer's shortage cost. On the other hand, when the supplier keeps the system inventory, the inventory level (S^s) is unaffected by an increase in the buyer's shortage cost. The optimal location of inventory depends on which cost is lower. Initially, suppose a system in which the buyer keeps the inventory, i.e., $l_d^* = l_p + 1$. Imagine that the shortage cost of the buyer is steadily increasing. The inventory level and associated costs are also rising. Eventually, the cost of all the added inventory at the buyer will increase the total system cost beyond the constant cost of having the supplier hold the inventory, and so the optimal inventory location will switch from the buyer to the supplier. Alternately, when p^b is low such that $\beta < 1$, the buyer keeps the system inventory. As p^b increases, eventually $\beta > 1$, i.e., the expected marginal cost of the buyer is larger than that of the supplier at optimality.[9] When this happens the supplier carries the system inventory. Relatively speaking, the same will happen should the supplier's expediting cost be significantly dropped (instead of raising the buyer's shortage cost).

Recall that the penalty p^s can also be considered to represent the supplier's flexibility. That is, a low value of p^s implies that the supplier can expedite at a lower cost and indicates high flexibility, and vice versa. Thus one implication of Proposition 2 is that in a supply chain a more flexible supplier is likely to also carry the system inventory. Flexibility to respond to changes in demand has been one of the features of the Just-in-Time phenomenon.

---

[9] It is straightforward to show that β is increasing in p^b.

Furthermore, as mentioned in the Introduction, anecdotal evidence suggests that a move to Just-in-Time systems also led to shifting of inventory to the suppliers. Our results suggest that this is indeed a system optimal behavior when the suppliers are flexible.

2.3.1 Effect of Transportation Lead Time

Most often, there is a finite non-zero transportation lead time for delivery of goods from a supplier to a buyer. This alters the model slightly, in that the supplier's cost must now reflect that the maximum amount of time it can take to get goods to the buyer has increased by l_t. That is, in (5) and hence in Lemma 1, l_p is replaced by $\bar{l}_p = l_p + l_t$. We can then use our model to interpret this new scenario in one of the following two ways.

1. Transportation lead time, l_t, could be included in the delivery lead time, l_d, as a lower bound; i.e., $l_d \geq l_t$. Instead of comparing the costs at $l_d = 0$ and $l_d = l_p + 1$ as in (8), we compare the costs between $l_d = l_t$ and $l_d = \bar{l}_p + 1$. Proposition 1 holds with the following modifications. First β is modified to $\bar{\beta}$ as follows:

   $$\bar{\beta} = \frac{\beta\sqrt{l_p + l_t + 1}}{\sqrt{l_p + 1} + \beta\sqrt{l_t}}.$$

   Second, the buyer always takes some risk and holds at least enough inventory to cover the demand during the transportation time. That is, the two extreme solutions include either the buyer carrying all of the inventory or the buyer carrying enough inventory to buffer against the transportation lead time and the supplier carrying the rest. Proposition 2, similarly, still holds as well, though one must also fix l_t and l_p.

2. Alternatively, transportation lead time, l_t, could be incorporated into the supplier's production lead time. We would then compare the costs of $l_d = 0$ with those of $l_d = \bar{l}_p + 1$. Here, β and hence Propositions 1 and 2 would remain unchanged. In this case, $l_d = 0$ would be equivalent to assuming that the supplier has a depot immediately next to the buyer and the production lead time is counted from the time the order is placed until it is produced and delivered to the depot. In some cases, it may be that the supplier actually keeps inventory in the buyer's warehouse, retail store or shop floor, as the case may be. This is the well known consignment situation or VMI; similar consignment

arrangements were derived by Lee & Whang (1999). This is a scenario widely practiced by corporations such as IBM Rochester, General Motors (Saturn), Walmart, and Guardian Insulation. Guardian Insulation, in cooperation with certain hardware suppliers, such as ACE Hardware, offers retailers the option of having the insulation on hand, but not actually paying for it until it is sold. It remains on the inventory list of Guardian (through ACE) and is not counted as inventory for the retailer. Guardian does, however, select the amount of inventory to keep on site at the retailer.

3 A Single-Supplier Multiple-Buyers Model

We now consider a model with a single supplier and multiple buyers. For this scenario, we assume that the supplier may negotiate a different delivery lead time with each of the buyers. We also assume that the demand distributions of the buyers are Normal and independent of each other, but not necessarily identical. Finally, we assume that costs are identical for all buyers. For non-identical costs, the notation becomes summarily complex and interpretation or manipulation difficult, at best.

3.1 Two Buyers

We will first develop the cost function for the problem with only two buyers. Let l_{d_i} denote the lead time to be negotiated with buyer i for $i = 1, 2$. We still assume that the supplier follows an order-up-to policy as before. She faces a normally distributed demand with mean

$$\mu_s = (l_p - l_{d_1} + 1)\mu_1 + (l_p - l_{d_2} + 1)\mu_2$$

and standard deviation

$$\sigma_s = \sqrt{(l_p - l_{d_1} + 1)\sigma_1^2 + (l_p - l_{d_2} + 1)\sigma_2^2}.$$

Thus we assume, as before, that the supplier raises her on-hand and relevant pipeline inventory to a base stock level, $S^s = \mu_s + k^s \sigma_s$. We must, however, modify what we call the supplier's relevant pipeline inventory, since she has two buyers with possibly differing delivery lead times. Without loss of generality, let us assume that $l_{d_1} \leq l_{d_2}$. As before, we can reduce the supplier's system to one with an effective production lead time of $l_p - l_{d_1}$ instead of

l_p. Moreover, in the new system, l_{d_1} will equal zero, and l_{d_2} will equal the difference between the original delivery lead times. To see what we consider pipeline inventory, suppose that the supplier can label all of her inventory according to which buyer generated it. The pipeline inventory will then be all of the demand generated by the first buyer and the demand of the second buyer which has been withdrawn from inventory but whose generated order has not yet finished production.

Let us turn now to the costs of operating this system.

Lemma 2 *The expected periodic cost function for the system with one supplier and two buyers, with identical costs and normally distributed independent demands, can be written as*

$$\begin{aligned} C(l_{d_1}, l_{d_2}) = (h^b + p^b)\phi(k^b)\left(\sigma_1\sqrt{l_{d_1}} + \sigma_2\sqrt{l_{d_2}}\right) \\ + (h^s + p^s)\phi(k^s)\sqrt{(l_p - l_{d_1} + 1)\sigma_1^2 + (l_p - l_{d_2} + 1)\sigma_2^2} \end{aligned} \tag{10}$$

$C(l_{d_1}, l_{d_2})$ *is concave in* (l_{d_1}, l_{d_2}) *with a minimum occurring at one of the extrema of the function where* $l_{d_1} = l_{d_2}$.

Thus the minimum for $C(l_{d_1}, l_{d_2})$ will occur where $l_{d_i} \in \{0, l_p + 1\}$. We then need to compare

$$C(0,0) = (h^s + p^s)\phi(k^s)\sqrt{\sigma_1^2 + \sigma_2^2}\sqrt{l_p + 1)} \tag{11}$$

$$C(l_p + 1, l_p + 1) = (h^b + p^b)\phi(k^b)(\sigma_1 + \sigma_2)\sqrt{l_p + 1} \tag{12}$$

The ratio between the above two cost functions is

$$\frac{C(l_p + 1, l_p + 1)}{C(0,0)} = \beta\frac{\sigma_1 + \sigma_2}{\sqrt{\sigma_1^2 + \sigma_2^2}}. \tag{13}$$

Note that, in contrast to (8) for the single-buyer single-supplier system, the standard deviations of the buyers' demand distributions now influence the decision of who should carry the inventory. This is a consequence of the risk pooling effect of multiple buyers. Proposition 3 below follows directly from (13).

Proposition 3 *Let* β *be as defined in Proposition 1. If* $\beta \leq \frac{\sqrt{\sigma_1^2+\sigma_2^2}}{\sigma_1+\sigma_2}$ *then it is optimal for the buyers to keep all of the system's inventory, and for the supplier to operate a make-to-order system. Otherwise, it is optimal for the supplier to keep the inventory for both buyers, and for the buyers to hold zero inventory.*

Observe that if $\beta = 1$ (that is, the expected marginal holding and penalty costs of the buyer and supplier are identical), then the supplier will hold all the inventory. This is the classical risk pooling argument. Thus Proposition 3 is a generalization of the classical risk pooling result to the case when costs at the buyer and supplier differ. When the expected marginal cost of the supplier sufficiently exceeds that of the buyers it offsets any benefit one would get from risk pooling. Consequently, the system is better off with buyers carrying their own inventory.

The concave cost function indicates that an extreme solution will be optimal, with a buyer keeping either all or none of the inventory. The above Proposition further restricts the solution space by ruling out the "split inventory" cases where only one of the buyers keeps inventory, and the supplier keeps the remaining system inventory. It is clear that such a solution cannot be optimal by arguing as follows. Such a scenario is equivalent to two systems, each composed of a single supplier and buyer, where the suppliers are identical and the buyers differ only in their demand distributions. Since the optimal location of each system's inventory depends only on the cost parameters (see Lemma 1) it must be true that in both, system inventory is optimally kept in the same location and never "split". Thus in a system with two buyers with identical costs either the two buyers keep the inventory or neither of them keeps inventory.

Furthermore, the scenarios where the buyers will carry their inventory have reduced in number as compared to the single-buyer single-supplier case. For instance, if $\sigma_1 = \sigma_2$, then only if $\beta < .7071$ will the buyers keep their own inventory, as opposed to $\beta < 1$ in the case of only one buyer. Hence, *the presence of multiple buyers makes it more likely that the supplier should offer zero delivery lead time.* This, however, depends on how different the two standard deviations are. For example, when the standard deviation of demands that the two buyers face are farther apart, then it is more likely that the buyers will keep their own inventory. Formally,

Corollary 1 *Let β be as defined in Proposition 1. Suppose that β and σ_i, $i \in \{1, 2\}$ are such that $\beta > \frac{\sqrt{\sigma_1^2+\sigma_2^2}}{\sigma_1+\sigma_2}$. By Proposition 3 the supplier carries inventory for both the buyers. Then as σ_i increases for a fixed β and σ_{-i}, there exists a threshold σ_i^* such that, for $\sigma_i > \sigma_i^*$ the buyers keep their own inventories.*

The result follows from Proposition 3 and the fact that $\frac{\sqrt{\sigma_1^2+\sigma_2^2}}{\sigma_1+\sigma_2}$ is increasing in σ_i for a fixed σ_{-i}.

Finally observe that it is optimal to offer identical delivery lead times to both buyers. This provides support for the identical delivery lead time assumption of Inderfurth (1994). Unfortunately, we will shortly see that the result does not necessarily generalize to more than two buyers.

3.2 Three or More Buyers

We now generalize these results to the case of $M > 2$ buyers with identical costs. Lemma 3 below is an extension of Lemma 2 for M buyers.

Lemma 3 *The expected periodic cost function for the system with one supplier and M buyers with identical costs and normally distributed demands can be written as*

$$C(l_{d_1}, \ldots, l_{d_M}) = (h^b + p^b)\phi(k^b)\left(\sum_{i=1}^{M} \sigma_i \sqrt{l_{d_i}}\right) + (h^s + p^s)\phi(k^s)\sqrt{\sum_{i=1}^{M} \sigma_i^2 (l_p - l_{d_i} + 1)} \tag{14}$$

$C(l_{d_1}, \ldots, l_{d_M})$ is concave in $(l_{d_1}, \ldots, l_{d_M})$ with a minimum occurring at one of the extrema of the function where $l_{d_i} \in \{0, l_p + 1\}\ \forall\, i = 1, \ldots, M$.

In Proposition 3 we saw that in a system with two buyers, either the supplier keeps inventory for both the buyers or the buyers hold their respective inventory. This extreme point solution, however, does not generalize to the case with more than two buyers, as we see next.

Proposition 4 *Let β be as defined in Proposition 1. Consider $M \geq 3$ buyers with $\sigma_1 \leq \sigma_2 \leq \cdots \leq \sigma_M$.*

1. *If $\beta \geq 1$, then the supplier keeps the system inventory for all buyers.*
2. *If $\beta < \beta_{min}$, then all buyers keep their own inventory, where*

$$\beta_{min} \equiv \min_k \left\{ \frac{\sqrt{\sum_{i=1}^{k} \sigma_i^2}}{\sum_{i=1}^{k} \sigma_i} \right\}.$$

3. *Otherwise (if $\beta_{min} \leq \beta < 1$), then there exists a buyer $k \geq 2$ such that the optimal system inventory allocation will be for the supplier to hold inventory for buyers 1 though k, and for buyers $k+1$ through M to hold their own inventory. The threshold k satisfies the following conditions:*

(a) $\sigma_k \leq \frac{2\beta}{1-\beta^2}\sqrt{\sum_{i=1}^{k-1}\sigma_i^2}$

(b) $\sigma_{k+1} \geq \frac{2\beta}{1-\beta^2}\sqrt{\sum_{i=1}^{k}\sigma_i^2}$

The first part of Proposition 4 implies that the supplier's expected marginal cost is lower than that of the buyers and hence she keeps all the inventory. This is the classical risk pooling argument. The second part can be considered a generalization of a similar condition for the case of two buyers stated in Proposition 3. The third part represents the situation when some buyers keep their own inventory while the supplier keeps inventory for others. Specifically, in determining which buyers will hold their own inventory, we need to only look among partitions of buyers that split the rank ordered buyers by the magnitude of their standard deviations.

Moreover, it is interesting to note that when the supplier keeps the inventory, it is for buyers with the *smallest* standard deviations. This is perhaps counter-intuitive, as one might expect that risk-pooling benefits are larger when the inventory of buyers with larger standard deviations are combined. This can be explained as follows. There are two effects that need to be considered. The first is the effect of standard deviation of demands on the benefit due to risk pooling. Observe that the benefit of risk pooling increases as the sets of standard deviations being pooled are "closer" to each other. Specifically, it is maximized when buyers have identical standard deviations. On the other hand, the benefit decreases as the standard deviations become "farther apart" (also see Corollary 1. The other factor to consider is the cost structure of the buyer and the supplier. Recall that β is the ratio of the expected marginal holding and penalty cost of the buyer to that of the supplier at optimality. Observe that when $\beta = 1$, the supplier keeps inventory for all the buyers. This is the classical risk pooling argument. The benefit may be large or small depending on the range of standard deviations being pooled. Regardless, there is always a positive benefit from pooling under this scenario. When $\beta < 1$, however, the expected marginal costs for the supplier is larger than the corresponding cost for the buyer. So by decreasing β from 1 to $1 - \epsilon$, $\epsilon > 0$, it may be optimal to shift some inventory from the supplier to one or more buyers. Which buyer should this be? Since the risk pooling benefits are more for sets of standard deviations closer to each other, the buyers with standard deviations on the extreme ends (low or high) should be carrying their own inventories. With $\beta < 1$ (with the interpretation above), the system gains most if the buyer with the largest standard deviation is shifted.

The following proposition illustrates how the optimal inventory allocation gets affected by change in system parameters.

Proposition 5 *Consider a system with one supplier and* $M \geq 3$ *buyers where* $\sigma_1 \leq \sigma_2 \leq \cdots \leq \sigma_M$*. Assume that the M buyers are optimally partitioned into two groups,* $\mathcal{S}$ *and* $\mathcal{B}$ *such that the supplier carries the inventory of all buyers in group* $\mathcal{S}$ *and buyers in group* $\mathcal{B}$ *carry their own inventory; that is* $\mathcal{S} = \{1, \ldots, k\}$ *and* $\mathcal{B} = \{k+1, \ldots, M\}$*.*

Consider the four cost parameters in the model – the holding and penalty costs for the supplier (h^s *and* p^s*) and for the buyer (*h^b *and* p^b*). For any given values of three of these parameters, there exist (i) two unique threshold value functions of the fourth parameter, (ii) an* m *such that* $1 \leq m \leq |\mathcal{S}|$ *but* $m \neq 2$*, and (iii) an* n*, such that* $|\mathcal{S}| + 1 \leq n \leq M$*, which redefine the optimal partition of buyers as follows:*

Parameters		*Optimal Partitioning & Conditions*		
Fixed	*Thresholds*	$\mathcal{S} \leftarrow \mathcal{S} \setminus \{m, \ldots, k\}$ $\mathcal{B} \leftarrow \mathcal{B} \cup \{m, \ldots, k\}$	*No Change*	$\mathcal{S} \leftarrow \mathcal{S} \cup \{k+1, \ldots, n\}$ $\mathcal{B} \leftarrow \mathcal{B} \setminus \{k+1, \ldots, n\}$
h^s, p^s, *and* h^b	$p_l^{b*}(k), p_h^{b*}(k)$	$p^b < p_l^{b*}(k)$	$p^b \in [p_l^{b*}(k), p_h^{b*}(k)]$	$p^b > p_h^{b*}(k)$
h^s, p^s, *and* p^b	$h_l^{b*}(k), h_h^{b*}(k)$	$h^b < h_l^{b*}(k)$	$h^b \in [h_l^{b*}(k), h_h^{b*}(k)]$	$h^b > h_h^{b*}(k)$
h^b, p^b, *and* h^s	$p_l^{s*}(k), p_h^{s*}(k)$	$p^s > p_h^{s*}(k)$	$p^s \in [p_l^{s*}(k), p_h^{s*}(k)]$	$p^s < p_l^{s*}(k)$
h^b, p^b, *and* p^s	$h_l^{s*}(k), h_h^{s*}(k)$	$h^s > h_h^{s*}(k)$	$h^s \in [h_l^{s*}(k), h_h^{s*}(k)]$	$h^s < h_l^{s*}(k)$

The above proposition is similar to Proposition 2 for a single-supplier and single-buyer case with analogous interpretations.

Suppose that the supplier is presently working with a set of buyers with the optimal location of inventory as determined above. Now a new buyer approaches the supplier to consider a similar inventory placement arrangement. What delivery lead time should the supplier quote this new buyer? How does it affect its relationship with the existing set of buyers? The following proposition provides an answer.

Proposition 6 *Consider a system with one supplier and M buyers divided into two groups,* $\mathcal{S}$ *and* $\mathcal{B}$*. If it is optimal for the supplier to keep the system inventory for the group of buyers in* $\mathcal{S}$ *while the buyers in* $\mathcal{B}$ *keep their own inventory, then adding another buyer to the system will result in exactly one of the following scenarios being optimal.*

1. *Buyer $M+1$ is added to group $\mathcal{S}$ while group $\mathcal{B}$ remains unchanged.*

2. *Buyer $M+1$ is added to group $\mathcal{B}$ while group $\mathcal{S}$ remains unchanged.*

3. *Buyer $M+1$ is added to group $\mathcal{S}$ along with q of the buyers with the smallest standard deviations in group $\mathcal{B}$, for some q.*

From Proposition 4 we know that group $\mathcal{S}$, say of size $m \leq M$ contains the buyers with the m smallest standard deviations. Hence Proposition 6 says the following.

1. If the new buyer has a standard deviation *smaller* than σ_m, then this buyer will be added to group $\mathcal{S}$, possibly along with some of the buyers from group $\mathcal{B}$. This is illustrated in the following example. Let $\sigma_1 = 1$, $\sigma_2 = 3$, $\sigma_3 = 5$, $\sigma_4 = 20$, $\sigma_5 = 80$, and $\beta = .7$. Then the minimum cost occurs when $\mathcal{S} = \{\sigma_1, \sigma_2, \sigma_3\}$.
   (a) Adding a sixth buyer with $\sigma_6 = 4$ results in a new minimal partition where $\mathcal{S} = \{\sigma_1, \sigma_2, \sigma_6, \sigma_3\}$.
   (b) Adding a sixth buyer with $\sigma_6 = 4.5$ results in a new minimal partition where $\mathcal{S} = \{\sigma_1, \sigma_2, \sigma_6, \sigma_3, \sigma_4\}$.

2. If the new buyer has a standard deviation *larger* than σ_m, then this buyer will either be added to $\mathcal{B}$ by himself, or added to $\mathcal{S}$ along with several other buyers who previously held their own inventory. Clearly, should he be added to $\mathcal{S}$, then at least all buyers with smaller standard deviations will also be included in $\mathcal{S}$ along with $M+1$. These two cases can be seen in the following examples.
   (a) Assume a buyer with $\sigma_6 = 250$ is added to the buyers previously listed. The new minimum occurs where $\mathcal{S}$ remains unchanged but where $\mathcal{B} = \{\sigma_4, \sigma_5, \sigma_6\}$.
   (b) Assume a buyer with $\sigma_6 = 21$ is added. In this case the new minimum occurs where $\mathcal{S} = \{\sigma_1, \sigma_2, \sigma_3, \sigma_4, \sigma_5, \sigma_6\}$ and where $\mathcal{B} = \emptyset$.

In other words, it will never happen that adding a buyer to an existing system will result in some of the buyers, for whom the supplier had previously carried inventory, now carrying their own inventory. This is an important "stability" result in the sense that contracting with more buyers does not upset the existing arrangements with present buyers in a negative fashion.

3.3 Effect of Transportation Lead Times

Once again, we can incorporate transportation lead times into our model. As before the supplier's cost must now reflect that the maximum amount of time it can take to get goods to the buyers has increased by l_t.[10] That is, in Lemma 4, l_p is replaced by $l_p + l_t$. Again, the model can be used to investigate two different policies. When transportation lead time is included in the production lead time, (comparing costs when $l_{d_i} = 0$ and $l_{d_i} = l_p + l_t + 1$), the propositions and lemmas remain unchanged. When it is included in the delivery lead time, (we examine $l_{d_i} = l_t$ instead of $l_{d_i} = 0$), the results still hold, though some conditions are slightly altered. In Proposition 4, β is replaced with $\beta \left(\sqrt{1 + \frac{l_t}{l_p+1}} - \sqrt{\frac{l_t}{l_p+1}} \right)$ and in Proposition 5, we must add l_t and l_p to the "fixed" column. Proposition 6 holds without any change. With both policies, the discussion remains the same as with only one buyer, and so is not repeated here.

4 Conclusion

This paper has examined delivery lead time and its impact on the costs of a two-stage serial and arborescent supply chain. We have demonstrated the link between delivery lead time and the location of system safety stock. For a system with only one buyer and one supplier, we showed that either the supplier or the buyer, but not both, should keep the system safety stock. Moreover, there is a threshold value for the production lead time that is a function of the remaining parameters, such that for production lead times larger (smaller) than this threshold, the supplier (buyer) should keep the system safety stock. Similar thresholds exist for both the supplier's expediting cost and the buyer's shortage cost.

For a system with one supplier and multiple buyers with identical cost parameters, we showed that the supplier will keep the safety stock for the buyers with the m smallest standard deviations, while the rest of the buyers kept their own inventory. This m is a function of the parameters of the system. We also showed that adding another buyer to an optimally

---

[10] In this section we assume identical transportations times. This can be relaxed. The analysis, however, becomes considerably complex and some of the results may alter. For instance, if we have two buyers with differing lead times, it is possible that one buyer holds inventory and the other not, so long as the buyer which holds the inventory has the shorter lead time and, if $l_{t_i} < l_{t_j}$, $\frac{\sigma_i^2}{\sigma_j^2} < 4\frac{l_{t_j}}{l_p+1}\left(1 + \frac{l_{t_j}}{l_p+1}\right)$.

operating system will never result in an optimal solution where a buyer, who previously had inventory held by the supplier, will now carry inventory.

In this work, we have operated from a centralized position, seeking to examine the optimal delivery lead time that minimizes total system cost. The optimal allocation of inventory and the changes wrought by adding a buyer to an existing system, however, raise interesting incentive issues, especially for the supplier. Unless appropriately compensated, she will likely be unwilling to shoulder the risk of the whole system by offering a delivery lead time of zero. Moreover, she may not be inclined to take an extra customer if re-setting the delivery lead times may cause her expenses to increase disproportionately to her revenue. Thus, a look at the incentive structure required to expand this work to the decentralized scenario is an area of future research.

5 Appendix

Proof of Lemma 1: The sum of (2) and (5) gives the expected system cost as follows:

$$C(l_d) = p^s(l_p - l_d + 1)\mu + p^b l_d \mu - (h^b + p^b)\int_0^{S^b} x f_{l_d}(x)dx - (h^s + p^s)\int_0^{S^s} x f_{l_p - l_d + 1}(x)dx \tag{15}$$

For a normal density $f(x)$, (15) can be rewritten as follows by observing that $\int_{-\infty}^{a} xf(x)dx = \mu F(a) - \sigma^2 f(a)$.

$$\begin{aligned}
C(l_d) &= \mu\left(p^b l_d + p^s(l_p - l_d + 1)\right) - (p^b + h^b)l_d\left(\mu F_{l_d}(S^b) - \sigma^2 f_{l_d}(S^b)\right) \\
&\quad - (p^s + h^s)(l_p - l_d + 1)\left(\mu F_{l_p - l_d + 1}(S^s) - \sigma^2 f_{l_p - l_d + 1}(S^s)\right) \\
&= \mu\left(p^b l_d + p^s(l_p - l_d + 1)\right) - (p^b + h^b)\left(\mu l_d \frac{p^b}{p^b + h^b} - \sigma\sqrt{l_d}\phi(k^b)\right) \\
&\quad - (p^s + h^s)\left(\mu(l_p - l_d + 1)\frac{p^s}{p^s + h^s} - \sigma\sqrt{l_p - l_d + 1}\phi(k^s)\right) \\
&= \sigma\left((h^s + p^s)\phi(k^s)\sqrt{l_p - l_d + 1} + (h^b + p^b)\phi(k^b)\sqrt{l_d}\right)
\end{aligned} \tag{16}$$

Taking the first and second derivatives of $C(l_d)$ with respect to l_d, we get:

$$\begin{aligned}
\frac{\partial C(l_d)}{\partial l_d} &= .5\sigma\left(\frac{(p^b + h^b)\phi(k^b)}{\sqrt{l_d}} - \frac{(p^s + h^s)\phi(k^s)}{\sqrt{l_p - l_d + 1}}\right) \\
\frac{\partial^2 C(l_d)}{\partial {l_d}^2} &= -.25\sigma\left(\frac{(p^b + h^b)\phi(k^b)}{l_d^{1.5}} + \frac{(p^s + h^s)\phi(k^s)}{(l_p - l_d + 1)^{1.5}}\right)
\end{aligned} \tag{17}$$

Clearly, the second partial will always be negative and so $C(l_d)$ is strictly concave. Hence its minima are to be found at the extremal values that l_d can take. ■

Proof of Lemma 2: The cost for buyer $i, i = 1, 2$ has the same functional form as the single buyer of Lemma 1:

$$C^{b_i}(l_{d_1}, l_{d_2}) = p^b l_{d_i} \mu_i - (h^b + p^b) \int_0^{S^{b_i}} x f_{l_{d_i}}(x) dx \,. \tag{18}$$

Using the same reasoning as before, the expected periodic cost function for the supplier can be derived. Where $g(\cdot)$ is a normal density with mean $\mu_s = (l_p - l_{d_1} + 1)\mu_1 + (l_p - l_{d_2} + 1)\mu_2$ and standard deviation $\sigma_s = \sqrt{(l_p - l_{d_1} + 1)\sigma_1^2 + (l_p - l_{d_2} + 1)\sigma_2^2}$, the expected periodic cost function for the supplier which is the two buyer equivalent of (5) is

$$C^s(l_{d_1}, l_{d_2}) = p^s \mu_s - (h^s + p^s) \int_0^{S^s} x g(x) dx \tag{19}$$

Therefore, the total expected periodic cost is written as follows:

$$\begin{aligned} C(l_{d_1}, l_{d_2}) &= C^s(l_{d_1}, l_{d_2}) + C^{b_1}(l_{d_1}, l_{d_2}) + C^{b_2}(l_{d_1}, l_{d_2}) \\ &= p^s \mu_s - (h^s + p^s) \int_0^{S^s} x g(x) dx + p^b l_{d_1} \mu_1 - (h^b + p^b) \int_0^{S^{b_1}} x f_{l_{d_1}}(x) dx \\ &\quad + p^b l_{d_2} \mu_2 - (h^b + p^b) \int_0^{S^{b_2}} x f_{l_{d_2}}(x) dx \\ &= (h^b + p^b)\phi(k^b) \left(\sigma_1 \sqrt{l_{d_1}} + \sigma_2 \sqrt{l_{d_2}} \right) \\ &\qquad + (h^s + p^s)\phi(k^s) \sqrt{(l_p - l_{d_1} + 1)\sigma_1^2 + (l_p - l_{d_2} + 1)\sigma_2^2} \end{aligned} \tag{20}$$

To show that $C(l_{d_1}, l_{d_2})$ is also a concave function in both l_{d_1} and l_{d_2}, we need to find the first and second partials.

$$\begin{aligned} \frac{\partial C(l_{d_1}, l_{d_2})}{\partial l_{d_i}} &= -\frac{\sigma_i^2 (h^s + p^s)\phi(k^s)}{2\sqrt{(l_p - l_{d_1} + 1)\sigma_1^2 + (l_p - l_{d_2} + 1)\sigma_2^2}} + \frac{\sigma_i (h^b + p^b)\phi(k^b)}{2\sqrt{l_{d_i}}} \\ \frac{\partial^2 C(l_{d_1}, l_{d_2})}{\partial {l_{d_i}}^2} &= -\frac{\sigma_i^4 (h^s + p^s)\phi(k^s)}{4((l_p - l_{d_1} + 1)\sigma_1^2 + (l_p - l_{d_2} + 1)\sigma_2^2)^{1.5}} - \frac{\sigma_i (h^b + p^b)\phi(k^b)}{4(l_{d_i})^{1.5}} \end{aligned} \tag{21}$$

$$\frac{\partial^2 C(l_{d_1}, l_{d_2})}{\partial l_{d_1} \partial l_{d_2}} = -\frac{\sigma_1^2 \sigma_2^2 (h^s + p^s)\phi(k^s)}{4\Big((l_p - l_{d_1} + 1)\sigma_1^2 + (l_p - l_{d_2} + 1)\sigma_2^2\Big)^{1.5}} \tag{22}$$

It is quite easily shown that the Hessian is negative definite, and hence the cost function is strictly concave in l_{d_1} and l_{d_2}. Hence the minimum for $C(l_{d_1}, l_{d_2})$ will occur where $l_{d_i} \in \{0, l_p + 1\}$.

$$C(0, 0) = (h^s + p^s)\phi(k^s)\sqrt{(l_p + 1)(\sigma_1^2 + \sigma_2^2)} \tag{23}$$

$$C(0, l_p + 1) = (h^b + p^b)\phi(k^b)\sigma_2\sqrt{l_p + 1} + (h^s + p^s)\phi(k^s)\sigma_1\sqrt{l_p + 1} \tag{24}$$

$$C(l_p + 1, 0) = (h^b + p^b)\phi(k^b)\sigma_1\sqrt{l_p + 1} + (h^s + p^s)\phi(k^s)\sigma_2\sqrt{l_p + 1} \tag{25}$$

$$C(l_p + 1, l_p + 1) = (h^b + p^b)\phi(k^b)(\sigma_1 + \sigma_2)\sqrt{l_p + 1} \tag{26}$$

We will show that if $C(l_p+1, l_p+1) \geq C(0, l_p+1)$ then $C(0,0) \leq C(0, l_p+1)$, that is $l_{d_1} = 0$ and $l_{d_2} = l_p + 1$ will never be optimal. First, using the above four equalities, we get

$$C(l_p+1, l_p+1) \geq C(0, l_p+1) \Leftrightarrow (h^b + p^b)\phi(k^b) \geq (h^s + p^s)\phi(k^s) \tag{27}$$

$$\begin{aligned} C(0, l_p+1) \geq C(0,0) &\Leftrightarrow (h^b + p^b)\phi(k^b)\sigma_2 \geq (h^s + p^s)\phi(k^s)\left(\sqrt{\sigma_1^2 + \sigma_2^2} - \sigma_1\right) \\ &\Leftrightarrow (h^b + p^b)\phi(k^b)\frac{\sigma_2}{\sqrt{\sigma_1^2 + \sigma_2^2} - \sigma_1} \geq (h^s + p^s)\phi(k^s) \end{aligned} \tag{28}$$

Since $\sigma_1 + \sigma_2 \geq \sqrt{\sigma_1^2 + \sigma_2^2}$, then $\frac{\sigma_2}{\sqrt{\sigma_1^2+\sigma_2^2}-\sigma_1} \geq 1$, and the result follows. Thus $l_{d_1} = 0$ and $l_{d_2} = l_p + 1$ will never be optimal. The case for $C(l_p + 1, 0)$ is identical and hence is omitted. Thus it will never be optimal for $l_{d_1} \neq l_{d_2}$, and the optimum will occur at either $l_{d_1} = l_{d_2} = 0$ or $l_{d_1} = l_{d_2} = l_p + 1$. ■

Proof of Proposition 4: To make the presentation of the proof easier, let us define some initial notation. Let $\mathcal{A}$ be a subset of buyers in the set of all buyers $\mathcal{S}$, and $\overline{\mathcal{A}}$ be its complement in $\mathcal{S}$, where $\mathcal{S}$ has cardinality M. Then for any $\mathcal{A}$ with cardinality m, let $C(m)_{\overline{\mathcal{A}}}$ be the total cost for the situation where the buyers in $\mathcal{A}$ keep their own inventory, and the supplier keeps all other buyers' inventory. Finally, recall that $\beta = \frac{(h^b+p^b)\phi(k^b)}{(h^s+p^s)\phi(k^s)}$ and define $\kappa = (h^s + p^s)\phi(k^s)\sqrt{l_p + 1}$. Then, the cost equations are

$$C(m)_{\overline{\mathcal{A}}} = \kappa\left(\beta\sum_{i\in\mathcal{A}}\sigma_i + \sqrt{\sum_{i\in\overline{\mathcal{A}}}\sigma_i^2}\right) \tag{29}$$

First, consider the case where $\beta \geq 1$.

$$\begin{aligned} C(m)_{\overline{\mathcal{A}}} = \kappa\left(\beta\sum_{i\in\mathcal{A}}\sigma_i + \sqrt{\sum_{i\in\overline{\mathcal{A}}}\sigma_i^2}\right) &\geq \kappa\left(\beta\sqrt{\sum_{i\in\mathcal{A}}\sigma_i^2} + \sqrt{\sum_{i\in\overline{\mathcal{A}}}\sigma_i^2}\right) \\ \geq \kappa\left(\sqrt{\sum_{i\in\mathcal{A}}\sigma_i^2} + \sqrt{\sum_{i\in\overline{\mathcal{A}}}\sigma_i^2}\right) &\geq \kappa\left(\sqrt{\sum_{i\in\mathcal{A}}\sigma_i^2 + \sum_{i\in\overline{\mathcal{A}}}\sigma_i^2}\right) = C(0)_{\mathcal{S}} \end{aligned} \tag{30}$$

The first and third inequalities come from properties of the square root and the second inequality comes from $\beta \geq 1$. Hence, the supplier keeps all of the system inventory when $\beta \geq 1$.

Now, consider the case where $\beta < 1$. First, we will show that either all buyers hold their own inventory, or there exists a buyer $k \geq 2$ such that the optimal system inventory allocation will be for the supplier to hold inventory for buyers 1 through k, and for buyers $k+1$ through M to hold their own inventory. Then we will demonstrate exactly when it will be optimal for all buyers to hold their own inventory.

To start, we will show that it will never be optimal for the supplier to hold inventory for only one buyer. Consider $\overline{\mathcal{A}} = \{k\}$ and $\mathcal{A} = \mathcal{S} \setminus \{k\}$. Then

$$C(M-1)_{\{k\}} = \kappa\left(\beta\sum_{i\in\mathcal{S}\setminus\{k\}}\sigma_i + \sigma_k\right) \geq \kappa\left(\beta\sum_{i\in\mathcal{S}\setminus\{k\}}\sigma_i + \beta\sigma_k\right) = \kappa\beta\sum_{i\in\mathcal{S}}\sigma_i = C(M)_{\emptyset} \tag{31}$$

Hence, when $\beta < 1$, the system will always prefer all the buyers holding their own inventory over the supplier holding inventory for only one buyer.

Next, for any $j \in \overline{\mathcal{A}} \neq \emptyset$, consider $C(m)_{\overline{\mathcal{A}}}$ and $C(m+1)_{\overline{\mathcal{A}}\setminus\{j\}}$.

$$\begin{aligned} C(m)_{\overline{\mathcal{A}}} \leq C(m+1)_{\overline{\mathcal{A}}\setminus\{j\}} &\iff \beta \sum_{i\in\mathcal{A}} \sigma_i + \sqrt{\sum_{i\in\overline{\mathcal{A}}} \sigma_i^2} \leq \beta \sum_{i\in\mathcal{A}\cup\{j\}} \sigma_i + \sqrt{\sum_{i\in\overline{\mathcal{A}}\setminus\{j\}} \sigma_i^2} \\ &\iff \sigma_j \leq \frac{2\beta}{1-\beta^2} \sqrt{\sum_{i\in\overline{\mathcal{A}}\setminus\{j\}} \sigma_i^2} \end{aligned} \tag{32}$$

For any $l \in \mathcal{A} \neq \emptyset$, consider $C(m)_{\overline{\mathcal{A}}}$ and $C(m-1)_{\overline{\mathcal{A}}\cup\{l\}}$.

$$\begin{aligned} C(m)_{\overline{\mathcal{A}}} \leq C(m-1)_{\overline{\mathcal{A}}\cup\{l\}} &\iff \beta \sum_{i\in\mathcal{A}} \sigma_i + \sqrt{\sum_{i\in\overline{\mathcal{A}}} \sigma_i^2} \leq \beta \sum_{i\in\mathcal{A}\setminus\{l\}} \sigma_i + \sqrt{\sum_{i\in\overline{\mathcal{A}}\cup\{l\}} \sigma_i^2} \\ &\iff \sigma_l \geq \frac{2\beta}{1-\beta^2} \sqrt{\sum_{i\in\overline{\mathcal{A}}} \sigma_i^2} \end{aligned} \tag{33}$$

Hence, if $C(m)_{\overline{\mathcal{A}}}$ were the minimum total cost, then $\sigma_j < \sigma_l$, for all $l \in \mathcal{A}$ and $j \in \overline{\mathcal{A}}$. Since the σ's are ordered by increasing size, this means that $C(m)_{\overline{\mathcal{A}}}$ can only be the minimum if $\overline{\mathcal{A}} = \{1, 2, \ldots, k\}$ and $\mathcal{A} = \{k+1, \ldots, M\}$ for some $k \in \{1, \ldots, M-1\}$ or if either $\overline{\mathcal{A}} = \emptyset$ or $\mathcal{A} = \emptyset$. Finally, $k \neq 1$ as the supplier will never hold inventory for only one buyer.

To show the necessary conditions for the minimum, first note that

$$\sigma_l \geq \frac{2\beta}{1-\beta^2} \sqrt{\sum_{i\in\overline{\mathcal{A}}} \sigma_i^2}$$

can be rewritten for some k as

$$\sigma_l \geq \frac{2\beta}{1-\beta^2} \sqrt{\sum_{i=1}^{k} \sigma_i^2}\ . \tag{34}$$

Since $\sigma_{k+1} \leq \sigma_{k+2} \leq \cdots \leq \sigma_M$, this condition, which must be true for all $l \in \mathcal{A}$, is equivalent to

$$\sigma_{k+1} \geq \frac{2\beta}{1-\beta^2} \sqrt{\sum_{i=1}^{k} \sigma_i^2}\ . \tag{35}$$

Similarly, we can find an appropriate k such that rewriting gives

$$\sigma_j \leq \frac{2\beta}{1-\beta^2} \sqrt{\sum_{i\in\overline{\mathcal{A}}\setminus\{j\}} \sigma_i^2} = \frac{2\beta}{1-\beta^2} \sqrt{\sum_{i=1}^{k} \sigma_i^2 - \sigma_j^2}\ . \tag{36}$$

¿From this we get that, for all $j \in \overline{\mathcal{A}}$

$$\sigma_j \leq \sigma_k \leq \frac{2\beta}{1-\beta^2} \sqrt{\sum_{i=1}^{k-1} \sigma_i^2} \leq \frac{2\beta}{1-\beta^2} \sqrt{\sum_{i=1}^{k-1} \sigma_i^2 + \sigma_k^2 - \sigma_j^2} \tag{37}$$

and hence the second condition which must be true for all $j \in \overline{\mathcal{A}}$ is equivalent to

$$\sigma_k \leq \frac{2\beta}{1-\beta^2}\sqrt{\sum_{i=1}^{k-1}\sigma_i^2}\ . \tag{38}$$

Finally, we will show that when $\beta < \beta_{\text{min}} = \min_k \left\{ \frac{\sqrt{\sum_{i=1}^{k}\sigma_i^2}}{\sum_{i=1}^{k}\sigma_i} \right\}$, the buyers will all keep their own inventory, which implies that the above derived conditions are the necessary conditions for the optimal partition when $\beta_{\text{min}} \leq \beta < 1$. From the definition of β_{min}, we can get that, for all $k \in \{1, 2, \ldots, M\}$

$$C(M)_{\emptyset} = \kappa\beta\sum_{i=1}^{M}\sigma_i < \kappa\sqrt{\sum_{i=1}^{k}\sigma_i^2} + \kappa\beta\sum_{i=k+1}^{M}\sigma_i = C(M-k)_{\overline{\mathcal{A}}} \tag{39}$$

where $\overline{\mathcal{A}} = \{1, 2, \ldots, k\}$. Hence, for $\beta < \beta_{\text{min}}$, it will be optimal for all buyers to hold their own inventory, and the theorem is proved. ■

Proof of Proposition 5: Given Proposition 4, let us introduce easier notation for the expected cost. Let C_k^M be the cost associated with a system of M buyers where k of them, those with the k smallest standard deviations, keep their inventory with the supplier. Then $C_k^M = \kappa\left(\beta\sum_{i=k+1}^{M}\sigma_i + \sqrt{\sum_{i=1}^{k}\sigma_i^2}\right)$, $C_0^M = \kappa\beta\sum_{i=1}^{M}\sigma_i$, and $C_M^M = \kappa\sqrt{\sum_{i=1}^{M}\sigma_i^2}$. Assume that C_m^M was the optimal cost for the case with M buyers. For all other values of k, this implies that

$$\beta\sum_{i=m+1}^{M}\sigma_i + \sqrt{\sum_{i=1}^{m}\sigma_i^2} \leq \beta\sum_{i=k+1}^{M}\sigma_i + \sqrt{\sum_{i=1}^{k}\sigma_i^2}.$$

It is easy to see that if β is increased, then C_k^M increases faster than C_m^M for all $k < m$. Hence, as β is increased, m will never decrease. On the other hand, if β is increased, then C_k^M increases slower than C_m^M for all $k > m$. Hence, it is possible that m could increase. It need not, however, increase smoothly. I.e., if β should increase by .1, then m could jump from 3 to 8.

Similarly, if β is decreased, then C_k^M decreases slower than C_m^M for all $k > m$. Hence, as β is decreased, m will never increase. On the other hand, if β is decreased, then C_k^M decreases faster than C_m^M for all $k > m$. Hence, it is possible that m could decrease. It need not, however, decrease smoothly. I.e., if β should decrease by .1, then m could jump from 3 to 0.

Finally, recall that $\beta = \frac{(h^b+p^b)\phi(k^b)}{(h^s+p^s)\phi(k^s)}$. This means that β is increasing in h^b and p^b, and decreasing in h^s and p^s, and the proposition is proved. ■

Proof of Proposition 6: Once again, let C_k^M be the cost associated with a system of M buyers where k of them, those with the k smallest standard deviations, keep their inventory with the supplier. Then

$C_k^M = \kappa\left(\beta \sum_{i=k+1}^M \sigma_i + \sqrt{\sum_{i=1}^k \sigma_i^2}\right)$. Assume that C_m^M was the optimal cost for the case with M buyers. For all other values of k, this implies that

$$\beta \sum_{i=m+1}^{M} \sigma_i + \sqrt{\sum_{i=1}^{m} \sigma_i^2} \leq \beta \sum_{i=k+1}^{M} \sigma_i + \sqrt{\sum_{i=1}^{k} \sigma_i^2}.$$

Assume we add a buyer with σ_{M+1}. Let q be the buyer such that $\sigma_q \leq \sigma_{M+1} \leq \sigma_{q+1}$. Then the minimal total cost function must have the following form.

$$C_k^{M+1} = \begin{cases} \kappa\left(\beta \sum_{i=k+1}^M \sigma_i + \beta\sigma_{M+1} + \sqrt{\sum_{i=1}^k \sigma_i^2}\right) & \text{for } k = 1, \ldots, q \\ \kappa\left(\beta \sum_{i=k+1}^M \sigma_i + \sqrt{\sum_{i=1}^k \sigma_i^2 + \sigma_{M+1}^2}\right) & \text{for } k = q+1, \ldots, M \end{cases} \tag{40}$$

Case 1: Assume $\sigma_{M+1} \geq \sigma_m$. This means that $m \leq q$. Hence, for all other $k \leq q$, since $C_m^M \leq C_k^M$ and $C_k^{M+1} = C_k^M + \kappa\beta\sigma_{M+1}$, it must be that $C_m^{M+1} \leq C_k^{M+1}$. This means that for $\sigma_{M+1} \geq \sigma_m$, either the added buyer will hold his own inventory, or he, along with all other buyers with smaller standard deviations, as well as possibly some with higher standard deviations, will have the supplier hold his inventory.

Case 2a: Next, assume $\sigma_k \leq \sigma_{M+1} < \sigma_m$. Following the same argument as above, $C_k^{M+1} = C_k^M + \kappa\beta\sigma_{M+1}$ implies that $C_k^M + \kappa\beta\sigma_{M+1} \geq C_m^M + \kappa\beta\sigma_{M+1}$. But $C_m^M + \kappa\beta\sigma_{M+1}$ is not even among the possible candidates for being optimal, since Proposition 2 indicates that it is dominated by some other optimum j. Hence C_k^{M+1} cannot possibly be optimal.

Case 2b: Finally, assume that $\sigma_{M+1} < \sigma_k \leq \sigma_m$. To see that $C_k^{M+1} \geq C_m^{M+1}$, first note that

$$\begin{aligned} C_k^{M+1} &= C_k^M + \kappa\sqrt{\sum_{i=1}^{k} \sigma_i^2 + \sigma_{M+1}^2} - \kappa\sqrt{\sum_{i=1}^{k} \sigma_i^2} \\ &\geq C_m^M + \kappa\sqrt{\sum_{i=1}^{k} \sigma_i^2 + \sigma_{M+1}^2} - \kappa\sqrt{\sum_{i=1}^{k} \sigma_i^2} \end{aligned} \tag{41}$$

which will in turn be greater than C_m^{M+1} if

$$\sqrt{\sum_{i=1}^{k} \sigma_i^2 + \sigma_{M+1}^2} + \sqrt{\sum_{i=1}^{m} \sigma_i^2} \geq \sqrt{\sum_{i=1}^{m} \sigma_i^2 + \sigma_{M+1}^2} + \sqrt{\sum_{i=1}^{k} \sigma_i^2}\ .$$

This will be true if and only if $m \geq k$, which is the case in this situation. Hence we have that $C_k^{M+1} \geq C_m^{M+1}$. Thus, for $\sigma_{M+1} \leq \sigma_m$, we get the same result. That is, it will never be optimal for an added buyer with standard deviation smaller than σ_m to join the system and hold his own inventory, thus causing several other buyers who had previously held their stock with the supplier to hold it themselves. He, possibly along with other buyers who had previously held inventory, will have the supplier hold it.

Thus, either the new buyer will be added to the group which keeps its own inventory without disturbing the location of the other buyers' inventory, or he will be added to the group for which the supplier keeps inventory, possibly along with some buyers who previously held their own inventory. ■

References

ARROW, K. J., SAMUEL KARLIN, AND HERBERT SCARF (1958). *Studies in the Mathematical Theory of Inventory and Production*, chapter 10. Stanford University Press.

CACHON, G. AND P. ZIPKIN (1999). Competitive and Cooperative Inventory Policies in a Two-Stage Supply Chain. *Management Science*, **45**(7), 936–953.

CHEN, F., ZVI DREZNER, JENNIFER K. RYAN, AND DAVID SIMCHI-LEVI (2000). Quantifying the Bullwhip Effect in a Simple Supply Chain: The Impact of Forecasting, Lead Times and Information. *Management Science*, **46**(3), 436–444.

CLARK, A. AND H. SCARF (1960). Optimal Policies for a Multi-Echelon Inventory Problem. *Management Science*, **6**, 475–490.

EPPEN, G. AND L. SCHRAGE (1981). *Centralized Ordering Policies in a Multi-Warehouse System with Lead Times and Random Demand*, (pp. 51–67). In *Studies in the Management Sciences, Vol. 16: Multi-Level Production/Inventory Control Systems: Theory and Practice*, L. Schwarz, ed., North-Holland Publishing Company.

FEDERGRUEN, A. AND P. ZIPKIN (1984). Approximations of Dynamic, Multilocation Production and Inventory Problems. *Management Science*, **24**(1), 69–84.

GALLEGO, G. AND P. ZIPKIN (1999). Stock Positioning and Performance Estimation in Serial Production-Transportation Systems. *Manufacturing & Service Operations Management*, **1**(1), 77–88.

GLASSERMAN, P. AND S. TAYUR (1995). Sensitivity Analysis for Base–Stock Levels in Multi–Echelon Production-Inventory Systems. *Management Science*, **41**(2), 263–281.

GRAVES, S. AND S. WILLEMS (1998). Optimizing Strategic Safety Stock Placement in Supply Chains. To appear in *Manufacturing & Service Operations Management*, **2**(1).

HENKOFF, R. (1994). Delivering the Goods. *Fortune*, **11**, 64–78.

INDERFURTH, K. (1994). Safety Stock Optimization in Multistage Inventory Systems. *International Journal of Production Economics*, **24**, 103–113.

LANGENHOFF, L. AND W.H.M. ZIJM (1990). An Analytical Theory of Multi-Echelon Production/Distribution Systems. *Statistica Neerlandica*, **44**(3), 149–172.

LEE, H. AND S. WHANG (1999). Decentralized Multi-Echelon Supply Chains: Incentives and Information. *Management Science*, **45**(5), 633–640.

LEE, H. L., V. PADMANABHAN, AND S. WHANG (1997). Information Distortion in a Supply Chain. *Management Science*, **43**(4), 546–558.

ROSLING, K. (1989). Optimal Inventory Policies for Assembly Systems under Random Demands. *Operations Research*, **37**, 565–579.

SIMPSON, K. (1958). In-Process Inventories. *Operations Research*, **6**, 863–873.