

Managing Supply Chain Demand Variability with Scheduled Ordering Policies

Gérard P. Cachon

Duke University, Fuqua School of Business, Durham, North Carolina 27708

www.duke.edu/~gpc

This paper studies supply chain demand variability in a model with one supplier and N retailers that face stochastic demand. Retailers implement scheduled ordering policies: Orders occur at fixed intervals and are equal to some multiple of a fixed batch size. A method is presented that exactly evaluates costs. Previous research demonstrates that the supplier's demand variance declines as the retailers' order intervals are balanced, i.e., the same number of retailers order each period. This research shows that the supplier's demand variance will (generally) decline as the retailers' order interval is lengthened or as their batch size is increased. Lower supplier demand variance can certainly lead to lower inventory at the supplier. This paper finds that reducing supplier demand variance with scheduled ordering policies can also lower total supply chain costs.

(*Supply Chain Management; Multi-Echelon Inventory; Bullwhip Effect*)

This paper studies the management of supply chain demand variability in a model with one supplier, N retailers, and stochastic demand. Retailers implement *scheduled ordering*: They may order only every T periods, and their order quantities must equal an integer multiple of a fixed batch size, Q_r .


Scheduled ordering policies influence the propagation of demand variance within a supply chain. Lee et al. (1997) demonstrate that the supplier's demand variance depends on the alignment of the retailers' orders. The supplier's demand variance is maximized when the retailers' orders are *synchronized*, i.e., all N retailers order in the same periods. It is minimized when the retailers' orders are *balanced*, i.e., the same number of retailers order each period. Assuming balanced orders, this paper demonstrates that the supplier's demand variance is further reduced when the retailer order intervals are lengthened (T is increased) or when the retailers' batch size is reduced.

The combination of these actions can dramatically dampen the supplier's demand variance.

To illustrate, Figure 1 plots two simulations of supplier demand with 16 retailers. Each retailer's mean demand equals one unit per period. The supplier's demand variance is clearly higher when the retailers may order in any period, $T = 1$. (In fact, when the retailers may order only every four periods, $T = 4$, the supplier's demand is no more variable than overall consumer demand.) Yet, in each case the retailers order, on average, every four periods.

In this model, lower supplier demand variance gives two primary benefits. One, for a fixed supplier fill rate, lower demand variance allows the supplier to carry less inventory on average. Two, for a fixed supplier average inventory, lower demand variance reduces the retailers' average lead time. (The retailers' lead time is the sum of two components, a fixed transportation time and a stochastic time due to

Figure 1 Simulated Supplier Demand and the Retailers' Ordering Policies: T = Order Interval Length; Q_r = Batch Size



inventory shortages at the supplier. It is this latter component that improves.)

However, reducing the supplier's demand variance is not free. Increasing T , *ceteris paribus*, raises a retailer's holding and backorder costs, because the retailer's ordering flexibility is reduced: A retailer may experience a large demand in a period, but may need to wait until some later period to order. Decreasing Q_r increases a retailer's order frequency, thereby raising ordering costs. It is also not clear that balancing the retailer order intervals will lower costs. With synchronized ordering the supplier can anticipate a large demand every T periods, so the supplier can arrange to have inventory arrive just before the retailers order.

This research has three objectives. First, demonstrate that scheduled ordering policies can reduce the supplier's demand variability. Second, present a method to evaluate supply chain costs exactly. Third, determine whether schedule ordering policies can reduce total supply chain costs.

In a numerical study it is found that balancing the retailers' orders does reduce costs. Lengthening the retailers' order interval, *ceteris paribus*, raises the supply chain's holding and backorder costs. A *flexible quantity* strategy is better: Lengthen the retailers' order interval and also reduce their batch size. This combination dramatically reduces the supplier's demand variance (as in Figure 1), continues to control the retailers' ordering costs, and can lower total supply chain costs. This strategy is effective when there are relatively few retailers and consumer demand vari-

ability is low. It is particularly effective if, in addition, the supplier is required to provide a high fill rate.

The next section summarizes the related literature. Section 2 details the model. Section 3 shows how scheduled ordering influences the supplier's demand variance. Section 4 evaluates scheduled ordering policies. Section 5 details the numerical study, and the final section concludes. All proofs are in the Appendix.

1. Literature Review

Lee et al. (1997) identify four causes of the *bullwhip effect*, the name given to the common observation that demand variance propagates up a supply chain. Synchronized ordering is one, as already mentioned. The other three are shortage gaming (retailers inflate their orders to receive a better allocation), demand updating (the supplier is unaware of true retailer demand and so must rationally assume a higher variance), and price fluctuations (retailers purchase more than their short term needs to take advantage of temporary price discounts). This paper does not consider those three causes. (Cachon and Lariviere (1996) study shortage gaming. Demand updating is studied by Drezner et al. (1996), and Chen et al. (1997).) Cohen and Baganha (1998) also study supply chain demand variance, but they do not consider strategies for reducing the variance of the retailers' orders.

In this paper there are five variables that influence the supplier's demand variance. Two are structural:

consumer demand variability and the number of retailers. The other three are policy parameters: the retailers' batch size, the retailers' order interval length, and the alignment of the retailers' order intervals.

Shapiro and Byrnes (1992) empirically examine demand variance in the medical supply industry (e.g. rubber gloves, saline solution, etc.). They observe that final demand exhibits little fluctuation, but orders from hospitals exhibit dramatic variability. As a remedy, they implemented standing order policies with the hospitals (a fixed amount is shipped on a regular schedule unless the customer specifically requests a different amount). The hospital required less space for storage, and the supplier's production efficiency improved. These results suggest that reducing the supplier's demand variance may benefit a supply chain.

Several papers concentrate on evaluating average costs. Deuermeyer and Schwarz (1981) and Svoronos and Zipkin (1988) provide techniques to approximate average costs in a continuous review model with Poisson demand. In similar settings, Axsater (1993) and Chen and Zheng (1997) provide exact methods. Cachon (1995) provides an exact algorithm for periodic review. Those papers assume no restrictions on when retailers may order. Chen and Samroengraja (1996) obtain exact results for a model in which retailers implement base stock policies ($Q_r = 1$) at fixed intervals and only one retailer orders at a time. This paper provides exact results with batch ordering, $Q_r > 1$, fixed order intervals, $T > 1$, and multiple retailers ordering in the same period. The technique extends Cachon (1995). A method is also presented to evaluate a nonstationary reorder point policy when retailer order intervals are synchronized.

Eppen and Schrage (1981) study a two-echelon model in which the supplier receives inventory at fixed intervals. The supplier carries no stock, so all inventory is immediately allocated among the retailers once it arrives at the supplier. Federgruen and Zipkin (1984), Jackson (1988), Jackson and Muckstadt (1989), McGavin et al. (1993), Nahmias and Smith (1994) and Graves (1996) allow shipments to retailers at intermediate times between replenishments to the supplier, thereby allowing the supplier to hold some stock. These models assume synchronized ordering (if a

shipment can occur to one retailer, then it can occur to any retailer) and unit ordering ($Q_r = 1$). The variability of supplier demand (i.e., retailer orders) has no impact, since the supplier is concerned only with the total amount of inventory needed at the start of each interval. This paper demonstrates that adjusting Q_r and T can hold the retailers' ordering frequency constant, yet improve supply chain performance by lowering the supplier's demand variance. There are also some significant structural differences: this paper allows the supplier to order each period, balanced order intervals, and batches ($Q_r > 1$).

Lee et al. (1996) and Aviv and Federgruen (1998) consider models in which retailers have fixed order intervals. They consider how information sharing can improve supply chain performance. Lee et al. (1996) assume synchronized ordering, and retailer orders are always filled either by the supplier or an outside source. Hence, the supplier's actions do not impact the retailers, nor do the retailers' actions influence the supplier's demand variance. Aviv and Federgruen (1998) consider both synchronized and balanced alignments. They find that balanced ordering generally has lower costs. They do not consider batch ordering nor do they study the supplier's demand variability. Their model is more complex than the one here, and their evaluations depend on approximations. (They have heterogeneous retailers and a supplier capacity constraint.)

Fixed interval ordering has been found to be very effective in multi-echelon models with deterministic demand (e.g. Roundy 1985, Maxwell and Muckstadt 1985). There, of course, supplier demand variance is not a relevant issue.

The quantity discount literature also concentrates on deterministic demand. Those models recommend that the supplier encourage retailers to increase their batch size. This advice is not necessarily appropriate when consumer demand is stochastic: Increasing Q_r will increase the supplier's demand variance.

2. Model

One supplier distributes a single product to N identical retailers. Within each period the following sequence of events occurs: (1) demand is realized; (2)

firms submit orders to their inventory sources; (3) shipments are released; (4) costs are assessed; and (5) shipments are received. Let the supplier be location 0 and the retailers locations 1 through N . In addition, identify the supplier with the subscript “ s ,” and a generic retailer with the subscript “ r .” Consumer demand is nonnegative, stationary, discrete, independent across retailers, and independent across periods. Let D_r^p be consumer demand over p periods, and let $\mu_r = E[D_r^1]$. For computational convenience, assume there exists a finite \bar{d} such that $\Pr\{D_r^1 \leq \bar{d}\} = 1$.

For each location i , the following are defined after demand occurs (event 1) but before orders are submitted (event 2): on-hand inventory, I_i ; backorders, B_i ; on-order inventory, OI_i ; and inventory position, $IP_i = I_i - B_i + OI_i$. (On-order inventory is inventory ordered but not received.) Let IP_i^+ and OI_i^+ be the inventory position and on-order inventory just after the firms order (event 2).

The retailers can order only in *review periods*, which occur every T periods. Retailers follow a scheduled ordering policy: In any review period when $IP_r \leq R_r$, a retailer orders a sufficient integer multiple of Q_r units to raise IP_r above R_r . Define a *batch* to be Q_r units. The supplier follows an (R_s, nQ_s) policy, which is analogous to the retailers’ policy, except the supplier may order in any period. Since the supplier’s demand equals an integer number of batches, all supplier variables are measured in batches of inventory, e.g. $I_s = 2$ means the supplier has $2Q_r$ units of inventory.

The supplier’s orders are always received in L_s periods. Inventory shipped from the supplier in period t arrives at a retailer in period $t + L_r$. Unfilled demands are backordered, and all backorders are eventually filled.

Since the supplier may receive orders from several retailers within a period, the supplier must allocate inventory among retailers. For analytical convenience, assume the supplier randomly shuffles the retailers’ orders for that period. This shuffling is independent of the retailer identities, order quantities, and shufflings from previous periods. So no retailer is given a preference. Once shuffled, orders are placed into an “order queue.” Orders are filled from this queue on a first-

in-first-out basis. Hence, if an order from period t is not filled in period t , this order will be filled before any order submitted in period $t + 1$ or later. Finally, the supplier will partially ship a retailer order.

Lee et al. (1997) divide scheduled ordering policies into three types, depending on how the retailers’ orders are aligned. When an equal number of retailers order per period, ordering is *balanced*. *Synchronized* ordering occurs when all N retailers order in the same periods, every T periods. Finally, *random* ordering occurs when orders are neither balanced nor synchronized.

Assume $T \leq N$. (If $N < T$, redefine period lengths so that $N = T$.) Balanced orders are possible only when N is an integer multiple of T . So, let $m^* = N/T$, and assume that m^* is an integer. Hence, m^* retailers order each period with balanced order intervals. (It is not difficult to extend the evaluation of policies to address non-integer m^* , but notational complexity is increased.)

There is a cost h_r per unit of retailer inventory per period, a cost h_s per unit of supplier inventory per period, and a cost p_r per retailer backorder per period. Let C be average supply chain costs per period,

$$C = N(h_r E[I_r] + p_r E[B_r]) + h_s Q_r E[I_s]. \quad (1)$$

Holding costs for on-route inventory are ignored.

3. A Retailer’s Ordering Processes

Lee et al. (1997) show that switching from synchronized to balanced retailer orders reduces the supplier’s demand variance, holding T constant and assuming $Q_r = 1$. This section assumes balanced ordering and investigates how changing T or Q_r affects the retailers’ ordering frequency and the supplier’s demand variance.

3.1. A Retailer’s Ordering Frequency

To control their ordering costs, the retailers must control their ordering frequency, ρ_r .

Theorem 1. *A retailer’s order frequency, ρ_r , declines as Q_r increases.*

There is a subtle relationship between ρ_r and T . As T increases, there are fewer review periods, but in

each review period there is a higher probability the retailer will submit an order. The former is stronger than the latter.

Theorem 2. *A retailer's order frequency, ρ_r , is nonincreasing in T . When $\Pr(D_r^{T+1} \geq Q_r + 1) > 0$, ρ_r decreases in T .*

3.2. Supplier's Demand Variance

From Theorem 1, increasing Q_r will reduce the retailer's ordering costs. But increasing Q_r also raises the supplier's demand variability.

Theorem 3. *The supplier's demand variance increases when the retailers' batch size, Q_r , is increased to jQ_r , where $j \in \{1, 2, 3 \dots\}$.*

Increasing T may also reduce the supplier's demand variance.

Theorem 4. *Assuming balanced ordering, increasing T lowers the supplier's demand variance when $Q_r > 1$. When $Q_r = 1$, the supplier's demand variance is independent of T .*

In addition to balancing retailer order intervals, these results suggest two strategies for managing the supplier's demand variance. The first strategy just increases T . This will reduce the retailers' ordering costs as well as the supplier's demand variance. Alternatively, a *flexible quantity* strategy increases T and decreases Q_r , thereby giving the retailer more flexibility to choose the order quantity but less flexibility in the timing of orders. Done correctly, this holds the retailers' ordering frequency relatively constant, thereby leaving ordering costs unchanged. Further, it will dramatically reduce the supplier's demand variance, as observed in Figure 1.

4. Evaluating Policies

This section presents a method to evaluate supply chain costs, assuming the retailers' order intervals are balanced. This assumption means that the supplier's demand process is stationary. Cachon (1995) provides exact results when $T = 1$. This method extends his approach to $T > 1$. All results are exact, unless otherwise noted.

For any R_s , this method evaluates the retailers' lead

time distribution: When the supplier has sufficient stock a retailer receives a batch in L_r periods, otherwise it is received with a larger delay. This distribution is used to evaluate other values of interest.

Section 4.7 addresses the evaluation of synchronized ordering. In that case the supplier's demand process is nonstationary, so a nonstationary reorder point policy is discussed.

4.1. Retailer Lead Time

To evaluate the retailer's lead time distribution, consider an arbitrary batch ordered by some retailer and track its progress through the supply chain. Averaging over all possible journeys through the supply chain yields the lead time distribution.

Begin with some definitions. Suppose retailer i has a review in period zero and submits an order. Given that an order was submitted, retailer i must have observed that his inventory position after demand in period zero was at or below the reorder point, $IP_r \leq R_r$. Define the overshoot random variable, $O_r = R_r - IP_r$. Hence, in period zero retailer i orders $\beta_r(O_r)$ batches,

$$\beta_r(o) = 1 + \lfloor o/Q_r \rfloor. \quad (2)$$

Consider the j th batch in this order, $j \in [1, \beta_r(o)]$. Define U_{oj} as the number of periods the supplier delays shipping the j th batch, conditional on $O_r = o$. (All variables with a subscript "o" are conditional on retailer i 's period zero overshoot.)

Cachon (1995) demonstrates that when the supplier implements a reorder point policy:

$$\Pr(U_{oj} \leq u) = \frac{1}{Q_s} \sum_{v=1}^{Q_s} \Pr(U_{ojv} \leq u), \quad (3)$$

where

$$\Pr(U_{ojv} \leq u) = \begin{cases} \Pr(XB_o^{L_s-u} \leq R_s + v - j), \\ R_s + v - j \geq 0, & 0 \leq u \leq L_s, \\ 1, & R_s + v - j \geq 0, \quad u \geq L_s + 1, \\ \Pr(XF_o^{u-L_s-1} > -1 - \beta_r(o) - (R_s + v - j)), \\ R_s + v - j < 0, & u \geq L_s + 1. \end{cases} \quad (4)$$

In the above, XB_o^p is the number of batches the retailers order over periods $[-p, 0]$, including only batches in

period zero placed in the supplier's order queue *before* retailer i 's order. XF_o^p is the number of batches the retailers order over periods $[0, p]$, including only batches in period zero placed in the supplier's order queue *after* retailer i 's order. Think of XB_o^p as the supplier's "time backwards" demand process and XF_o^p as the supplier's "time forward" demand process, both relative to retailer i 's order. Note that the above results are independent of T . The evaluations of XB_o^p and XF_o^p depend on T .

4.2. Supplier Demand Processes

The supplier's demand processes, XB_o^p and XF_o^p , are each divided into two components: Batches ordered by retailer i and batches ordered by the $N - 1$ "non- i " retailers. Let XNB^p be the number of batches ordered by the non- i retailers over periods $[-p, 0]$, including only batches in period zero placed in the supplier's order queue *before* retailer i 's order. Let XNF^p be the number of batches ordered by those retailers over periods $[0, p]$, including only batches in period zero placed in the supplier's order queue *after* retailer i 's order. Both XNB^p and XNF^p are independent of O_r because independent consumer demand implies independent retailer ordering processes.

Define YB_o^p as the number of batches retailer i orders over periods $[-p, -1]$, and define YF_o^p as the number of batches retailer i orders over periods $[1, p]$. Since retailer i 's ordering process is independent of the ordering process of the non- i retailers,

$$XB_o^p = YB_o^p + XNB^p, \quad p \geq 0, \tag{5}$$

and

$$XF_o^p = YF_o^p + XNF^p, \quad p \geq 0. \tag{6}$$

See the Appendix for the evaluation of YB_o^p and YF_o^p . Cachon (1995) determines that XNB^p and XNF^p have the same distribution. Therefore, for notational convenience, define XN^p as a random variable with the same distribution as XNB^p and XNF^p .

Before evaluating XN^p , some preliminary results and definitions are useful. Define Y_m^t as the number of batches m retailers order over t review periods. In steady state a retailer's inventory position after he orders is uniformly distributed on the interval $[R_r + 1, R_r + Q_r]$. Hence,

$$\Pr(Y_1^t \leq b) = \frac{1}{Q_r} \sum_{k=0}^{Q_r-1} \Pr(D_r^{tT} \leq bQ_r + Q_r - 1 - k). \tag{7}$$

The retailers' ordering processes are independent, so simple convolution yields

$$Y_m^t = Y_{m-1}^t + Y_1^t. \tag{8}$$

Now evaluate XN^p . Recognize that XN^p can be divided into two components: the batches ordered by retailers with a review in period zero and those ordered by retailers that don't have a review in period zero. There are $m^* - 1$ non- i retailers that have a review in period zero. Let \hat{Y}_n^p be the number of batches $n - 1$ non- i retailers order over periods $[-p, 0]$, including only those batches ordered by the retailers placed in the supplier's order queue before retailer i in period zero. That is, \hat{Y}_m^p is the first component of XN^p . Since retailer orders are shuffled each period, there is a $1/m^*$ probability that retailer i is the m th retailer order in period zero (including retailers that "order" zero batches). The $m - 1$ retailers before retailer i in the supplier's order queue will have $\tau(p)$ reviews included in \hat{Y}_m^p , where

$$\tau(p) = \lfloor p/T \rfloor + 1. \tag{9}$$

The $m^* - m$ retailers after retailer i in the supplier's order queue will have $\tau(p) - 1$ reviews included in \hat{Y}_m^p . Hence,

$$\Pr(\hat{Y}_m^p \leq b) = \frac{1}{m^*} \sum_{m=1}^{m^*} \Pr(Y_{m-1}^{\tau(p)} + Y_{m^*-m}^{\tau(p)-1} \leq b). \tag{10}$$

Overall,

$$XN^p = \hat{Y}_{m^*}^p + \sum_{j=1}^{\min\{T-1,p\}} Y_m^{\tau(p)-j}, \tag{11}$$

where the summation gives the batches ordered by the non- i retailers that don't have a review in period zero, i.e., the second component of XN^p .

4.3. Retailer Average Inventory Level

The supplier's demand processes, XB_o^p and XF_o^p , are used to evaluate the probability that the supplier delays shipping retailer i 's j th batch by u periods, $\Pr(U_{oj} = u)$. This delay is now used to evaluate a

retailer's average inventory, $E[I_r]$. These results are independent of T , so they are discussed briefly.

From Little's Law $E[I_r] = \mu_r E[S]$, where $E[S]$ is the expected sojourn for a unit of inventory (i.e., number of periods a unit is recorded in inventory). To evaluate $E[S]$, consider the c th unit in the j th batch of retailer i 's period zero order. If this unit is demanded in period p , the sojourn for this unit equals

$$S_{ojuc} = (p - u - L_r - 1)^+, \quad (12)$$

where u is the number of periods the supplier delays shipping the j th batch.

From Cachon (1995),

$$E[S_{ojuc}] = \sum_{p=w+L_r+1}^{\infty} \Pr(D_r^p \leq \phi_{ojc} - 1), \quad (13)$$

where

$$\phi_{ojc} = R_r - o + (j - 1)Q_r + c. \quad (14)$$

When a retailer's lead time demand is independent of the lead time, D_r^p is independent of $u + L_r$. In that case (13) is exact. This holds whenever $R_s \geq -1$. When $R_s < -1$, a retailer's lead time depends on the lead time demand, but this relationship is weak, especially when there are many retailers. Hence, when $R_s < -1$, (13) is an approximation.

Cachon (1995) demonstrates that $E[S_{ojuc}]$ can be evaluated with finite effort. Deconditioning across overshoots, batches, lead times and units yields $E[S]$,

$$E[S] = \frac{\frac{1}{Q_r} \sum_{o=0}^{\bar{o}} \sum_{j=1}^{\beta_r(o)} \sum_{u=0}^{\bar{u}} \sum_{c=1}^{Q_r} E[S_{ojuc}] \Pr(O_r = o) \Pr(U_{oj} = u)}{\sum_{o=0}^{\bar{o}} \Pr(O_r = o) \beta_r(o)} \quad (15)$$

where $\bar{o} = \bar{d}T - 1$ is the maximum overshoot and \bar{u} is the maximum shipping delay. (When $R_s \geq -1$, $\bar{u} = L_s + 1$, otherwise $\bar{u} \geq L_s + 1$.) See the Appendix for the evaluation of $\Pr(O_r = o)$.

4.4. Retailer Backorder Level

From $E[I_r]$ it is possible to evaluate $E[B_r]$. From the definition of a retailer's inventory position,

$$E[IP_r^+] = E[I_r] - E[B_r] + E[OI_r^+]. \quad (16)$$

IP_r^+ is uniformly distributed on the interval $[R_r + 1, R_r + Q_r]$, which on average equals $R_r + (Q_r + 1)/2$. At the start of a review period a retailer's inventory position is on average $R_r + (Q_r + 1)/2 - \mu_r(T - 1)$. Thus, the retailer's average inventory position is

$$E[IP_r^+] = R_r + \frac{1}{2}(Q_r + 1) - \frac{1}{2}\mu_r(T - 1). \quad (17)$$

According to Little's Law, $E[OI_r] = \mu_r(E[U] + L_r + 1)$, where $E[U]$ is the supplier's expected delay to ship a batch,

$$E[U] = \sum_{o=0}^{\bar{o}} \sum_{j=1}^{\beta_r(o)} E[U_{oj}] \Pr(O_r = o) / \sum_{o=0}^{\bar{o}} \Pr(O_r = o) \beta_r(o). \quad (18)$$

Hence,

$$E[B_r] = E[I_r] - R_r - \frac{1}{2}(Q_r + 1) + \mu_r(\frac{1}{2}(T - 1) + E[U] + L_r + 1). \quad (19)$$

4.5. Supplier Inventory and Fill Rate

Analogous to the evaluation of the retailer's average backorders, Cachon (1995) demonstrates that

$$E[I_w] = R_w + \frac{1}{2}(Q_w + 1) + \mu_r N(E[U] - L_s - 1) / Q_r. \quad (20)$$

$E[I_w]$ depends on T only through the evaluation of $E[U]$.

The supplier's average fill rate (percentage of batches shipped without delay) is $E[F_s]$,

$$E[F_s] = \sum_{o=0}^{\bar{o}} \sum_{j=1}^{\beta_r(o)} \Pr(U_{oj} = 0) \Pr(O_r = o) / \sum_{o=0}^{\bar{o}} \Pr(O_r = o) \beta_r(o). \quad (21)$$

4.6. Choosing Policies

For a given Q_r and T , reorder points are chosen to minimize total supply chain holding and backorder

costs, since the reorder points don't influence the retailers' ordering costs. C is convex in R_r , but not necessarily jointly convex in R_r and R_s . Therefore, a search is needed to find the optimal reorder points. See Axsater (1993) for additional details.

4.7. Synchronized Ordering

With synchronized ordering the supplier's demand is nonstationary. Nevertheless, the supplier could still choose to implement a stationary reorder point policy. Simplicity is a stationary policy's primary advantage. Alternatively, the supplier could try to improve performance with a nonstationary policy. A particular nonstationary policy is described below.

Assume the supplier chooses a stationary reorder point policy. The results in §4.1 continue to apply since they don't depend on the timing of the retailer orders. The evaluation of XN^p does change slightly, because all of the retailers order in the same period as retailer i . More specifically, (11) is simplified to

$$XN^p = \hat{Y}_N^p \tag{22}$$

All of the other evaluations continue to apply.

The supplier may carry more inventory than needed with a stationary policy. Suppose the supplier implements a stationary policy with reorder point R_s , retailers have reviews in periods $\{0, T, 2T, \dots\}$, and $L_s < T$. In period 0, the supplier's inventory position may fall to R_s or lower. In that case the supplier will order some inventory that arrives in period L_s . Some of this inventory may be used to fill backorders. But the rest of that inventory just sits at the supplier until period T , since there are no retailer orders over periods $L_s + 1$ to $T - 1$. Clearly, the supplier would have been better off delaying the arrival of some of that inventory until the end of period $T - 1$.

To formalize the above intuition, first assume $L_s < T$. (The alternative case could be handled, but with significantly greater analytical complexity.) Now let the supplier operate with two reorder points: R_s is applied in periods $\{[T - L_s - 1, T - 1], [2T - L_s - 1, 2T - 1], \dots\}$; and \hat{R}_s is applied in the other periods, $\{[0, T - L_s - 2], [T, 2T - L_s - 2], \dots\}$. R_s ensures that sufficient inventory arrives at the supplier just before the retailers order. \hat{R}_s has two purposes: ensure that the supplier handles backorders in

the same manner that she would if she operated with R_s as a single reorder point; and avoid receiving inventory before it could possibly be needed. Therefore, $\hat{R}_s = \min\{R_s, -1\}$: Replenishments to fill backorders are ordered in the same periods as they would be if the supplier operated with R_s in every period, but orders for inventory that the retailers could only request in the subsequent review period are delayed.

The retailers notice no difference between a supplier that operates with R_s as her single reorder point and a supplier that operates with the dual reorder points, (R_s, \hat{R}_s) . Hence, evaluation of $E[I_r]$ and $E[B_r]$ is the same as if the supplier operates with just R_s . Evaluation of the supplier's average inventory changes.

The supplier's average inventory equals her average inventory over one review cycle, periods $[0, T - 1]$. (Inventory is measured when costs are assessed.) Assume $R_s > -1$ (otherwise $\hat{R}_s = R_s$). After ordering in period $-L_s - 1$, the supplier's inventory position, IP_s , is uniformly distributed on the interval $[R_s + 1, R_s + Q_s]$. Further, there will be at most one outstanding order (because $L_s < T$). Hence, after the retailers order in period zero the supplier's inventory equals $(IP_s - Y_N^1)^+$.

If $IP_s - Y_N^1 \geq 0$, this inventory level will persist until the next review period. If $IP_s - Y_N^1 \leq -1$, the supplier immediately orders a sufficient number of batches to cover the backorders. The supplier will have zero inventory over periods $[0, L_s]$ (when inventory charges are assessed) and over periods $[L_s + 1, T - 1]$ the supplier may have some positive inventory (if $Q_s > 1$, the supplier may need to order more inventory than is needed just to cover the backorders in period zero). So the supplier's expected inventory over these T periods is

$$E[I_s] = E[(IP_s - Y_N^1)^+] + \frac{T - L_s - 1}{T} \times E \left[\left[\frac{(IP_s - Y_N^1)^- + Q_s - 1}{Q_s} \right] \cdot Q_s - (IP_s - Y_N^1)^- \right] \tag{23}$$

The term in the second expectation is the supplier's inventory in periods $[L_s + 1, T - 1]$ due to the supplier's period T order.

5. Numerical Study

A numerical study assesses the impact of scheduled ordering policies on supply chain performance. The following are held constant throughout the study: $L_s = L_r = h_s = h_r = \mu_r = 1$. The primary 48 problems are constructed from all combinations of the following parameters:

$$N \in \{4, 16\}; \hat{Q}_s \in \{1, 4\},$$

$$p_r \in \{1, 5, 25, 50\}, \sigma_r \in \{0.21, 1, \sqrt{2}\}.$$

\hat{Q}_s is the number of periods of average demand the supplier's batch size can satisfy. Average demand per period is N , so the supplier's batch size is $N\hat{Q}_s$ units. Since Q_s is measured in batches, $Q_s = N\hat{Q}_s/Q_r$. The parameter σ_r is the standard deviation of the demand distribution. When $\sigma_r = 0.21$, consumer demand at each retailer has a "discrete" normal distribution:

$$\Pr(D_r^1 \leq 0) = 0.02275,$$

$$\Pr(D_r^1 \leq 1) = 0.97725,$$

$$\Pr(D_r^1 \leq 2) = 1.$$

When $\sigma_r = 1$, the consumer demand distribution is Poisson, truncated so that $D_r^1 \leq 7$. When $\sigma_r = \sqrt{2}$, the consumer demand distribution is negative binomial,

$$\Pr(D_r^1 = d) = (1/2)^{d+1},$$

truncated so that $D_r^1 \leq 13$. These three distributions are chosen to represent situations with "low," "medium," and "high" demand variability. Agrawal and Smith (1994) find that the negative binomial distribution is appropriate in many retail environments.

The parameters T and Q_r have yet to be included. For each problem, several scenarios are created, where each scenario is one of the following combinations of T and Q_r

$$T \in \{1, 2, 4, \dots, N\},$$

$$Q_r \in \{1, 2, \dots, \min\{N\hat{Q}_s, 16\}\}.$$

Table 1 Change in Holding and Backorder Costs When Switching from Synchronized Order Intervals to Balanced Order Intervals

Demand	σ/μ	N	Minimum	Average	Maximum
Normal	0.21	4	-32.6%	-9.7%	0.0%
		16	-24.5%	-9.1%	0.0%
Poisson	1.00	4	-15.3%	-5.3%	0.0%
		16	-17.9%	-6.5%	0.0%
Negative Binomial	1.41	4	-11.0%	-3.3%	0.0%
		16	-15.0%	-5.1%	0.5%

The restriction that $T \leq N$ ensures that at least one retailer orders each period with balanced order intervals. The bound on the retailer's batch size, $Q_r \leq N\hat{Q}_s$, ensures that a retailer's minimum order quantity is no greater than the supplier's minimum order quantity, each measured in units. Overall, there are 216 scenarios.

5.1. Results

Table 1 presents data on the change in supply chain holding and backorder costs when the retailers' order interval alignment is switched from synchronized to balanced. With synchronized ordering the supplier implements the nonstationary reorder point policies discussed in §4.7. Only scenarios with $T > 1$ are considered, since there is no difference between the two alignments when $T = 1$.

The data are clear. Balancing order intervals significantly reduces supply chain holding and backorder costs. This strategy appears to be most effective with low consumer demand variability. This is somewhat surprising. If demand were deterministic, supply chain holding costs would be lower with synchronized ordering. (There would be no backorder costs since all demands are anticipated and the supplier could lower her holding costs.) However, with deterministic demand the supplier never risks running out of inventory, hence the retailers receive all shipments within L_r periods. Once some consumer demand variability is introduced, the supplier must carry safety stock to guarantee reliable deliveries. These data indicate that a little bit of uncertainty creates a strong incentive to smooth out the supplier's demand. The remaining discussion assumes balanced orders.

Table 2 Coefficient of Variation of a Single Retailer's Ordering Process

Demand Distribution	μ/σ	T	Q_r				
			1	2	4	8	16
Normal	0.21	1	0.21	1.00	1.73	2.65	3.87
		2	0.15	0.21	1.00	1.73	2.65
		4	0.11	0.15	0.21	1.00	1.73
		8	0.08	0.10	0.14	0.20	1.00
		16	0.05	0.07	0.09	0.13	0.18
Poisson	1.00	1	1.00	1.20	1.74	2.65	3.87
		2	0.71	0.79	1.07	1.73	2.65
		4	0.50	0.53	0.64	1.02	1.73
		8	0.35	0.36	0.40	0.53	1.00
		16	0.25	0.25	0.27	0.32	0.45
Negative Binomial	1.41	1	1.41	1.53	1.88	2.66	3.87
		2	1.00	1.05	1.24	1.76	2.65
		4	0.71	0.73	0.81	1.09	1.73
		8	0.50	0.51	0.54	0.65	1.02
		16	0.35	0.36	0.37	0.41	0.53

- Switching from synchronized to balanced order intervals reduces supply chain holding and backorder costs.

Increasing the length of the retailer's order interval is another strategy to reduce the supplier's demand variance. Table 2 displays data on the variability of a single retailer's order process. In all cases the supplier's demand variance will decline as T is increased. (The supplier's demand variance is N times a retailer's order variance.) However, increasing T always increased supply chain holding and backorder costs in these data. (In a broader experimental design, a few cases were found in which increasing T lowered supply chain costs a bit.) Further, Table 3 indicates that the average increase in costs is substantial. Hence, when considering supply chain holding and backorder costs, merely increasing T is not an appropriate strategy to reduce the supplier's demand variance.

- Lengthening the retailers' order intervals alone reduces the supplier's demand variability but increases supply chain holding and backorder costs.

Increasing T also reduces the retailer's ordering frequency, which will reduce ordering costs. But this benefit is not captured in Table 3. A flexible quantity

Table 3 Average % Increase in Holding and Backorder Costs Relative to $T = 1$

Demand Distribution	σ/μ	N	T			
			2	4	8	16
Normal	0.21	4	11%	36%		
		16	11%	40%	106%	248%
Poisson	1.00	4	7%	21%		
		16	8%	24%	60%	135%
Negative Binomial	1.41	4	7%	19%		
		16	7%	22%	51%	111%

Table 4 Retailer Ordering Frequency

Demand Distribution	μ/σ	T	Q_r				
			1	2	4	8	16
Normal	0.21	1	0.9772	0.5000	0.2500	0.1250	0.0625
		2	0.4997	0.4889	0.2500	0.1250	0.0625
		4	0.2500	0.2500	0.2447	0.1250	0.0625
		8	0.1250	0.1250	0.1250	0.1226	0.0625
		16	0.0625	0.0625	0.0625	0.0625	0.0614
Poisson	1.00	1	0.6321	0.4482	0.2489	0.1250	0.0625
		2	0.4323	0.3647	0.2406	0.1250	0.0625
		4	0.2454	0.2363	0.2012	0.1239	0.0625
		8	0.1250	0.1248	0.1231	0.1076	0.0624
		16	0.0625	0.0625	0.0625	0.0624	0.0563
Negative Binomial	1.41	1	0.5000	0.3750	0.2344	0.1245	0.0625
		2	0.3750	0.3125	0.2188	0.1235	0.0625
		4	0.2344	0.2188	0.1816	0.1190	0.0625
		8	0.1245	0.1235	0.1190	0.1005	0.0618
		16	0.0625	0.0625	0.0625	0.0618	0.0538

strategy attempts to lengthen T and reduce Q_r so as to keep the retailer's ordering frequency relatively constant. For example, the retailers could swap the values of Q_r and T , assuming $Q_r > T$. Table 4 indicates that any of those swaps will have only a small impact on the retailer's ordering frequency, but Table 2 indicates that the supplier's demand variance will decline substantially. Further reductions in the supplier's demand variance are possible if the retailers swap the value of Q_r and T , and then set $Q_r = 1$. However, Table 2 indicates that the retailer's order frequency may rise slightly. For example, with high demand variability,

Table 5 Change in Holding and Backorder Costs When a Flexible Quantity Strategy Is Adopted (T Is Increased and Q_r Is Decreased, but the Retailers' Order Frequency is Held Constant.)

Demand Distribution	σ/μ	N	Minimum	Average	Maximum
Normal	0.21	4	-28%	-13%	-2%
		16	-23%	-7%	6%
Poisson	1.00	4	-2%	5%	11%
		16	-7%	12%	32%
Negative Binomial	1.41	4	4%	9%	16%
		16	-2%	18%	45%

switching from ($Q_r = 4, T = 2$) to ($Q_r = 1, T = 4$) raises the ordering frequency from 0.2188 to 0.2344.

Table 5 presents data on the change in supply chain costs when a flexible quantity strategy is adopted. All of the possible swaps mentioned above are included in these data. On average, a flexible quantity strategy reduces supply chain costs when there are few retailers and low consumer demand variability. However, the strategy loses its effectiveness with an increase in either consumer demand variability or the number of retailers. Both of those effects can be explained. As consumer demand variability increases, Table 2 indicates that the flexible quantity strategy is less effective at reducing the retailers' order variance. As the number of retailers increases, the supplier's demand variance is reduced no matter their ordering policy. So as N increases, a reduction in the supplier's demand variance has a smaller impact. In fact, while a flexible quantity strategy always reduces the supplier's demand variance, it also imposes a cost on the retailers, i.e., it reduces their flexibility to time orders to demand surges and troughs. Even a small increase in each retailer's costs can, once totaled across many retailers, easily dwarf a significant decrease in the supplier's costs. (Looking only at minimum performance, it appears that the flexible quantity strategy is more effective with larger N . As N increases, more (Q_r, T) swaps are feasible. So that result is due to the larger number of observations.)

- A flexible quantity strategy reduces supply chain costs when there are few retailers and consumer demand variability is low.

Table 6 Change in Holding and Backorder Costs When a Flexible Quantity Strategy Is Adopted (T Is Increased and Q_r Is Decreased, but the Retailers' Order Frequency is Held Constant) and the Supplier Is Required to Provide a 99% Fill Rate

Demand Distribution	σ/μ	N	Minimum	Average	Maximum
Normal	0.21	4	-50%	-21%	9%
		16	-39%	-18%	14%
Poisson	1.00	4	-10%	3%	15%
		16	-28%	3%	24%
Negative Binomial	1.41	4	-3%	7%	16%
		16	-18%	11%	34%

Recall that one of the benefits of lower supplier demand variance is that the supplier can carry less inventory for a given fill rate. This benefit is likely to be strongest when the supplier is required to offer a high fill rate, say 99% or better. (In that case the second benefit of lower demand variance, improved retailer lead times, will be minimal since the retailers receive reliable deliveries in all cases.) In fact, there are many supply chains that operate under such a fill rate requirement; see Cachon and Fisher (1997) and Hart (1995). Note that this requirement raises in this setting overall supply chain holding and backorder costs, so it is assumed that it is adopted for reasons that are not explicitly modeled.

Table 6 presents data on supply chain costs when the supplier is required to offer a 99% fill rate and a flexible quantity strategy is adopted. (Since R_s must be an integer value, it is usually not possible to choose R_s to yield exactly a 99% fill rate. Therefore, the smallest and largest R_s are found that yield above and below 99%, respectively. Linear extrapolation of these two cases provides the cost estimate.) As in Table 5, the flexible quantity strategy is most effective when consumer demand variability is low and there are few retailers. Comparison between Tables 5 and 6 reveals that the flexible quantity strategy is more effective on average when the supplier is required to offer the 99% fill rate.

- Reducing the supplier's demand variance through a flexible quantity strategy is most effective when the supplier is required to offer a high fill rate.

6. Conclusion

An important lesson in supply chain management research is that firms should consider global supply chain performance, and not just the performance of their portion of the chain (see Lee and Billington 1992). *Ceteris paribus*, a reduction in the supplier's demand variance will reduce the supplier's average inventory. This research explores whether this also reduces total supply chain costs.

Two strategies were found to reduce the supplier's demand variance and also reduce total supply chain costs. The first, balancing retailer order intervals, is effective in a broad range of conditions. The second is a flexible quantity strategy: increase T and reduce Q_r so that the retailers' order frequency is held relatively constant. This strategy is effective when there are few retailers and consumer demand variability is low. It is particularly effective if, in addition, the supplier is required to provide a high fill rate.

This research highlights that the supplier's demand variance is an imperfect proxy for overall supply chain performance. For example, while increasing T will generally reduce the supplier's demand variance, that strategy also raises supply chain costs. Dampening the supplier's demand is only reasonable if (1) the supplier's costs represent a significant fraction of overall supply chain costs and (2) this action does not substantially raise the retailers' costs. Both of those conditions become less likely with increases in either N or the retailers' demand variance. Therefore, reducing a supplier's demand variance is an objective to adopt selectively.¹

¹ The author would like to thank Fangruo Chen, Marshall Fisher, Martin Lariviere, Hau Lee, Roy Shapiro, and Paul Zipkin for many helpful comments. The assistance of the Associate Editor and the anonymous referees is graciously acknowledged. Previous versions of this paper were titled "On the Value of Managing Supply Chain Demand Variability with Scheduled Ordering Policies."

Appendix A. Retailer i 's Ordering Processes and Overshoot

Given retailer i has a review in period zero, he will have $\tau(p) - 1$ reviews over periods $[1, p]$. Analogous to Y_m^t ,

$$\Pr(YF_o^p \leq b) = \Pr(D_r^{\tau(p)-1} \leq bQ_r - 1 - R_r + IP_r(o)), \quad p \geq 1, \tag{A-1}$$

where $IP_r(o)$ is retailer i 's inventory position at the start of period 1,

$$IP_r(o) = R_r + Q_r - o + Q_r \left\lfloor \frac{o}{Q_r} \right\rfloor. \tag{A-2}$$

Over periods $[-p, -1]$, retailer i also has $\tau(p) - 1$ reviews,

$$\Pr(YB_o^p \leq b) = \sum_{k=R_r+1}^{R_r+Q_r} \Pr(IP_r = k | o) \cdot \Pr(D_r^{\tau(p)-1} \leq bQ_r + R_r + Q_r - k), \quad p \geq 1, \tag{A-3}$$

where application of Bayes' theorem yields

$$\Pr(IP_r = k | o) = \frac{\Pr(D_r^T = o - R_r + IP_r)}{\Pr(D_r^T \leq Q_r + o) - \Pr(D_r^T \leq o)}. \tag{A-4}$$

At the end of a review period a retailer's inventory position, IP_r , is uniformly distributed on the interval $[R_r + 1, R_r + Q_r]$. Let d equal demand over periods $[1, T]$. Define $\hat{O} = R_r - (IP_r - d)$,

$$\Pr(\hat{O} = \hat{o}) = \frac{1}{Q_r} \sum_{j=0}^{Q_r-1} \Pr(D_r^T = Q_r + \hat{o} - j). \tag{A-5}$$

The retailer submits an order to the supplier in period T if $\hat{o} \geq 0$. In this case \hat{o} equals the retailer's overshoot. From Bayes' theorem, $\Pr(O_r = o) = \Pr(\hat{O} = o | o \geq 0)$.

Appendix B. Proofs

Proof of Theorem 1. The probability a retailer orders a positive quantity in a review is $1 - \Pr(Y_1^1 = 0)$, so

$$\rho_r = \frac{1 - \Pr(Y_1^1 = 0)}{T} = \frac{1}{T} \left(1 - \frac{1}{Q_r} \sum_{k=0}^{Q_r-1} \Pr(D_r^T \leq k) \right). \tag{B-1}$$

The result is immediate, since $\Pr(D_r^T \leq k)$ is nondecreasing in k . \square

Proof of Theorem 2. Say a request occurs in a period if the retailer would submit an order if $T = 1$. A request is called a *trigger request* if it is the first request since the last review period. Hence, once a trigger request occurs an order will certainly be submitted at the next review, no matter when future requests occur. Let $T = \alpha$, $\alpha \geq 1$. Show that the retailer's ordering frequency does not increase when T is increased to $\alpha + 1$.

Suppose a request occurs in period p . Define the random variable Δ such that the previous request occurred in period $p - \Delta$. Consider the first review to occur on or before period p . If period $p - \Delta$ is before this review, then the request occurring in period p is a trigger request. Since there is one trigger request for each order, the rate at which trigger requests occur equals the rate at which orders occur, i.e., the ordering frequency.

Divide time into consecutive groups of periods, where each group contains $\alpha(\alpha + 1)$ periods. Groups are chosen so that in the last

period of the group a review occurs whether the order interval is α or $\alpha + 1$. Let π equal the probability that a request is submitted in period p . Since demand is stationary, the probability that a request occurs in a period is independent of when the reviews occur. Therefore π is constant. Let $G(T, \delta)$ be the expected number of trigger requests per group, counting only those requests for which $\Delta = \delta$, and considering $T \in \{\alpha, \alpha + 1\}$,

$$G(T, \delta) = \pi(2\alpha + 1 - T) \min\{\delta, T\}. \tag{B-2}$$

When $\delta > \alpha$, $G(\alpha, \delta) = G(\alpha + 1, \delta)$, and when $\delta \leq \alpha$, $G(\alpha, \delta) > G(\alpha + 1, \delta)$. Hence, for all realizations of Δ , $G(\alpha, \delta) \geq G(\alpha + 1, \delta)$, which implies the ordering frequency is nonincreasing in T . If $\Pr(\Delta \leq \alpha) > 0$, then the ordering frequency will decrease in T (i.e., with positive probability, the second of two successive requests must occur within α periods of the first). This holds when $\Pr(D_r^{\alpha+1} \geq Q_r + 1) > 0$. \square

Proof of Theorem 3. It is sufficient to show that the variance of a single retailer's ordering process increases. Let $Y_{m,q}^t$ be the number of batches m retailers orders over t consecutive reviews, assuming these m retailers share the same review periods and q is their batch size. It holds that

$$V[Y_{1,q}^t] = q^2 E[(Y_{1,q}^t)^2] - E[(D_r^{tT})^2], \tag{B-3}$$

where $V[X]$ denotes the variance of the random variable X . Define $\omega_b = \Pr(Y_{1,q}^t = b)$. From Cachon (1995),

$$\omega_b = \frac{1}{q} \sum_{k=0}^{q-1} \Pr(D_r^{tT} \leq bq + q - 1 - k) - \Pr(D_r^{tT} \leq bq - 1 - k). \tag{B-4}$$

Hence,

$$q^2 E[(Y_{1,q}^t)^2] = q^2 \sum_{b=1}^{\infty} b^2 \omega_b. \tag{B-5}$$

Note that

$$(jq)^2 E[(Y_{1,qj}^t)^2] = q^2 \sum_{b=1}^{\infty} \left(\left[(2b - j) \left\lfloor \frac{b-1}{j} \right\rfloor - j \left\lfloor \frac{b-1}{j} \right\rfloor^2 + b \right] j \omega_b \right). \tag{B-6}$$

Since $E[(D_r^{tT})^2]$ is independent of q ,

$$V[Y_{1,qj}^t] - V[Y_{1,q}^t] = q^2 \sum_{b=1}^{\infty} \left(b - j \left\lfloor \frac{b}{j} \right\rfloor \right) \left(j \left\lfloor \frac{b}{j} \right\rfloor + j - b \right) \omega_b > 0. \tag{B-7}$$

\square

Proof of Theorem 4. Assume $N/T = m$, and m is an even integer, $m \geq 2$, so balanced ordering can be maintained even after doubling T . Before the order intervals are doubled, the supplier's demand variance per period is $mV[Y_r^1]$, and after doubling the order intervals, it is $mV[Y_r^2]/2$, where Y_r^2 is the order of a single retailer over two review periods, each of length T . It holds that $V[Y_r^2] = 2V[Y_r^1] + 2\text{Cov}(Y_r^1, Y_r^1)$. (The covariance is taken between two successive order intervals, each of length T .) It is clear that when $Q_r = 1$, the covariance of two consecutive orders is zero, so $V[Y_r^2] = 2V[Y_r^1]$. Hence, the supplier's demand variance is unchanged. When $Q_r > 1$, it can be shown that the covariance of two consecutive orders is negative, so $V[Y_r^2] < 2V[Y_r^1]$. \square

References

Agrawal, N., S. A. Smith. 1994. Estimating negative binomial demand for retail inventory management with lost sales. *Naval Res. Logistics* **43** 839–861.

Aviv, Y., A. Federgruen. 1998. The operational benefits of information sharing and vendor managed inventory (VMI) programs. Working Paper, Olin School of Business, Washington University, St. Louis, MO.

Axsater, S. 1993. Exact and approximate evaluation of batch-ordering policies for two-level inventory systems. *Oper. Res.* **41** 777–785.

Cachon, G. 1995. Exact evaluation of batch-ordering policies in two-echelon supply chains with periodic review. Working Paper, Fuqua School of Business, Duke University, Durham, NC.

—, M. Fisher. 1997. Campbell soup's continuous product replenishment program: evaluation and enhanced decision rules. *Production and Oper. Management* **6** 266–276.

—, M. Lariviere. 1996. Capacity choice and allocation: Strategic behavior and supply chain performance. Forthcoming in *Management Sci.*

Chen, F., J. Ryan, D. Simchi-Levi. 1997. The impact of exponential smoothing forecasts on the bullwhip effect. Working Paper, Northwestern University, Evanston, IL.

—, R. Samroengraja. 1996. A staggered ordering policy for one-warehouse multi-retailer systems. Forthcoming in *Oper. Res.*

—, Y.-S. Zheng. 1997. One-warehouse multi-retailer systems with centralized stock information. *Oper. Res.* **45** 275–287.

Cohen, M., M. Baganha. 1998. The stabilizing effect of inventory in supply chains. *Oper. Res.* **46** S72–S83.

Deuermeyer, B. L., L. B. Schwarz. 1981. A model for the analysis of system service level in warehouse-retailer distribution systems: the identical retailer case. *TIMS Studies in the Management Sciences* **16** 163–193.

Drezner, Z., J. Ryan, D. Simchi-Levi. 1996. Quantifying the bullwhip effect: the impact of forecasting, leadtime and information. Working Paper, Northwestern University, Evanston, IL.

Eppen, G., L. Schrage. 1981. Centralized ordering policies in a

- multi-warehouse system with lead times and random demands. *TIMS Studies in the Management Sciences* **16** 51–67.
- Federgruen, A., P. Zipkin. 1984. Computational issues in an infinite-horizon, multiechelon inventory model. *Oper. Res.* **32** 818–836.
- Graves, S. 1996. A multiechelon inventory model with fixed replenishment intervals. *Management Sci.* **42** 1–18.
- Hart, C. 1995. Internal guarantee. *Harvard Bus. Rev.* Jan–Feb. 64–71.
- Jackson, P. L. 1988. Stock allocation in a two-echelon distribution system or 'What to do until your ship comes in'. *Management Sci.* **34** 880–895.
- , J. A. Muckstadt. 1989. Risk pooling in a two-period, two-echelon inventory stocking and allocation problem. *Naval Res. Logistics* **36** 1–26.
- Lee, H., C. Billington. 1992. Managing supply chain inventory: pitfalls and opportunities. *Sloan Management Rev.* **33** 65–73.
- , P. Padmanabhan, S. Whang. 1997. Information distortion in a supply chain: the bullwhip effect. *Management Sci.* **43** 546–558.
- , K. C. So, C. Tang. 1996. Supply chain re-engineering through information sharing and replenishment coordination. Working Paper, Stanford University, Stanford, CA.
- Maxwell, W., J. Muckstadt. 1985. Establishing consistent and realistic reorder intervals in production-distribution systems. *Oper. Res.* **33** 1316–1341.
- 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.
- Nahmias, S., S. A. Smith. 1994. Optimizing inventory levels in a two echelon retailer system with partial lost sales. *Management Sci.* **40** 582–596.
- Roundy, R. 1985. 98%-effective integer ratio lot-sizing for one-warehouse multi-retailer systems. *Management Sci.* **31** 1416–1430.
- Shapiro, R., J. Byrnes. 1992. Intercompany operating ties: unlocking the value in channel restructuring. Working Paper, Harvard Business School, Cambridge, MA.
- Svoronos, A., P. Zipkin. 1988. Estimating the performance of multi-level inventory systems. *Oper. Res.* **36** 57–72.

Accepted by Hau Lee; received April 8, 1996. This paper has been with the author 11 months for 2 revisions.