EXACT EVALUATION OF BATCH-ORDERING INVENTORY POLICIES IN TWO-ECHELON SUPPLY CHAINS WITH PERIODIC REVIEW

GÉRARD P. CACHON

The Wharton School of Business, University of Pennsylvania, Philadelphia, Pennsylvania 19004
cachon@wharton.upenn.edu, opim.wharton.upenn.edu/~cachon

(Received May 1995; revisions received December 1995, August 1997, January 1999, September 1999; accepted October 1999)

This paper studies a two-echelon supply chain with stochastic and discrete consumer demand, batch order quantities, periodic inventory review, and deterministic transportation times. Reorder point policies manage inventories at every location. Average inventory, backorders and fill rates are evaluated exactly for each location. Safety stock is evaluated exactly at the lower echelon and a good approximation is detailed for the upper echelon. Numerical data are presented to demonstrate the model's utility. It is found that system costs generally increase substantially if the upper echelon is restricted to carry no inventory, of if the upper echelon is required to provide a high fill rate. In many cases it is optimal to set the upper echelon's reorder point to yield near zero safety stock, yet in some cases this simple heuristic can significantly increase supply chain operating costs. Finally, policies selected under the assumption of continuous inventory review can perform poorly if implemented in an environment with periodic review.

This paper studies a distribution system with one central warehouse and N identical retailers. Inventory is reviewed periodically and transportation times are deterministic. Firms implement reorder point policies and order quantities equal integer multiples of a fixed batch size. Consumer demand is stochastic with a known discrete distribution function that is stationary and independent across time and locations.

Average inventory, backorders and fill rates are evaluated *exactly* for each location in the system. Safety stock at the retail level is evaluated exactly, while a good approximation is given for safety stock at the warehouse.

This model is designed to reflect several features frequently observed in actual supply chains. For example, periodic review of inventory is common in practice. Significant flexibility is allowed in choosing the demand distribution, which is convenient since the best fitting distribution function varies by context. In spare parts applications the Poisson distribution accurately models demand (for a sample of studies, see Sherbrooke 1968, Muckstadt and Thomas 1980, Graves 1985, Cohen, Kleindorfer, and Lee 1986, and Hausman and Erkip 1994), but the negative binomial distribution can provide superior performance in retailing (see Nahmias and Smith 1994, and Aggrawal and Smith 1996). Flexibility is also allowed in choosing objectives. Some managers prefer to minimize inventory subject to meeting a minimum fill rate while others prefer to minimize total inventory holding and backorder costs. With either of those objectives optimal reorder point policies are found.

The studied model is related to several others in the multi-echelon literature with batching: Deurermeyer and Schwarz (1981), Moinzadeh and Lee (1986), Lee and Moinzadeh (1987a,b), Svoronos and Zipkin (1988); and Axsäter (1993). These studies also assume identical retailers, fixed transportation times, exogenously determined batch quantities, and independent and stationary consumer demand. However, they assume continuous review of inventory and Poisson consumer demand. The following use approximation to evaluate similar 2-echelon, N retailer models: Rosenbaum (1981), Schwarz, Deuermeyer, and Badinelli (1985), Lee and Billington (1993), Tempelmeier (1993), and Hausman and Erkip (1994). There are some exact results for models with continuous review. Axsäter (1993) provides exact results for Poisson demand and identical retailers, and Axsäter (2000) provides exact results for compound Poisson demand and nonidentical retailers. Chen and Zheng (1997) provide exact results when the central warehouse uses echelon stock reorder point policies. (Echelon stock policies are based on inventory information throughout the system, not just on inventory information at the central warehouse.) Cheung and Hausman (2000) provide exact results for the central warehouse assuming it serves nonidentical retailers. Liljenberg (1996) studies a model that is similar to this one with the exception that he assumes a different allocation policy at the central warehouse. (The allocation policy specifies the sequence in which the warehouse gives retailers inventory.)

In this paper an exact solution is provided for a system with batch ordering and periodic review. These features raise several complications. Batch ordering means that the demand process at the central warehouse is complex (i.e., it need not be a simple convolution of the retailers' demand processes). Periodic review means that a retailer

Subject classifications: Inventory/Production: multi-echelon, stochastic demand, periodic review heuristic.
Area of review: Logistics and Supply Chain Operations.

may order multiple batches within a period; hence, the batches within a single order might be shipped from the warehouse in different periods. Furthermore, the difference between a retailer's reorder point and its inventory position when it orders is stochastic, a source of variability that must be incorporated into an exact analysis. Finally, with periodic review the warehouse must adopt a policy for allocating inventory to retailers whenever total orders exceed available inventory. That allocation policy influences system performance; so, it too must be incorporated into an exact analysis.

The primary challenge in this exact solution is the evaluation of the retailer's lead time, which is the number of periods between when a retailer orders a batch and when it receives the batch from the central warehouse. A natural approach to evaluate the lead time distribution is to count the number of periods a retailer has to wait to receive each batch the warehouse orders. For example, say the warehouse orders a batch in period t that arrives at the warehouse at the end of period $t+L_w$ (it can only be shipped to a retailer in period $t+L_w+1$ or later). If a retailer orders this batch after period $t+L_w$, the warehouse can ship the batch immediately. However, if a retailer orders this batch in period $t+L_w$ or earlier, then the batch experiences a shipping delay. Svoronos and Zipkin (1988) suggested this technique for evaluating the lead time distribution, but unfortunately it is computationally burdensome, so they resorted to approximations.

A different perspective obtains computationally tractable exact results. Instead of evaluating when a batch is ordered by a retailer relative to the period the warehouse orders the batch (as described in the previous paragraph), evaluate when a batch is ordered by the warehouse relative to the period a retailer orders it from the warehouse. With both perspectives it is necessary to determine when a batch is ordered. However, the former must also evaluate *which retailer* orders it, whereas the latter must evaluate *which warehouse* orders it. In a system with one warehouse and N retailers it is clear that the "which warehouse" problem is much easier than the "which retailer" problem. Once the retailer's lead time distribution is evaluated, standard results yield the desired performance measures for each location in the supply chain, e.g., average inventory levels, backorders, and fill rates.

To demonstrate the utility of this model, numerical results are presented for 80 scenarios, 32 of which closely resemble scenarios studied by Svoronos and Zipkin (1988) and Axsäter (1993). Several observations are made from this sample. It is found, assuming warehouse holding costs equal retailer holding costs, that supply chain costs generally increase substantially if the warehouse reorder point is chosen to avoid holding any inventory at the warehouse, or, at the other extreme, if the warehouse reorder point is chosen to yield a high fill rate. Interestingly, many firms adopt the latter, "high warehouse fill rate" heuristic.

The warehouse reorder point can also be chosen so as to target a particular warehouse safety stock. One heuristic, suggested by Schwarz, Deuermeyer, and Badinelli (1985), sets the warehouse reorder point so that the warehouse safety stock equals $-Q_w$, where Q_w is the warehouse's batch size. Another heuristic sets the warehouse reorder point so that the warehouse safety stock equals zero. Graves (1996) found that a nonpositive warehouse safety stock is often optimal in a continuous review model with fixed interval shipments and one-for-one ordering. For the majority of scenarios tested those two heuristics perform reasonably well (i.e., supply chain costs within 10% of optimal costs), but there are some scenarios in which even those heuristics perform poorly, e.g., when high retailer fill rates are required (99% or higher). Also, it can be expected that their performance would decline if holding costs at the central warehouse were significantly lower than at the retailers.

A fifth heuristic implements the reorder points that are optimal for a continuous review model that provides an approximation for the actual periodic review model. That heuristic sometimes provides reasonable performance, in particular when consumer demand is low, but often yields costs that are significantly higher than optimal. Overall, it is concluded that it may be difficult to specify a simple heuristic for setting the warehouse's reorder point which provides good results in all settings.

The remainder of this paper is organized as follows. The next section outlines the assumptions that govern the operation of the system considered. The algorithm to evaluate performance measures for each location in the system is detailed in §2, and §3 presents the numerical study. The final section offers opportunities for future research.

1. THE MODEL

There is one central warehouse and N retail sites in this system. Let D^τ be demand at a single retailer over τ periods. Demand is measured in units, and all variables referring to the retailers are also measured in units. Demand is discrete, identically distributed across retailers, independent across retailers and time, and stationary across time. Further, there are two mild assumptions imposed on the demand distribution: D^1 is finitely bounded, i.e., there exists a $\bar{d}$ such that $\Pr(D^1 \le \bar{d}) = 1$, and $\Pr(D^1 = 1) > 0$.

Time is divided into periods of equal length. During a period the following events occur:

(1) demand occurs at each retailer;
(2) retailers request replenishment from the warehouse;
(3) the warehouse fills retailer orders and orders replenishments from its source;
(4) inventory and backorders are measured, and costs are charged; and
(5) the firms receive replenishments.

In each period the following variables are measured after demand (step 1) but before replenishments are requested (steps 2–3)

I on hand inventory;

B backorders;

IT on-order inventory (inventory ordered but not received);

IP inventory position, $IP = I - B + IT$.

A "w" subscript means the variable is associated with the warehouse, while an "r" subscript means the variable is associated with some retailer. A"$-$" superscript denotes the variable is measured at the beginning of a period (i.e., before demand), while a "$+$" superscript denotes the variable is measured at the end of period (i.e., after inventory arrives).

Replenishments for a retailer shipped from the warehouse in period t arrive at the retailer in period $t+L_r$. The warehouse's source has infinite capacity so the warehouse's replenishments requested in period t are always received in period $t+L_w$. All unfilled demands are backordered and eventually filled.

Retailer's orders are multiples of Q_r units, where the Q_r unit is called a *batch*. Retailers use an (R_r, nQ_r) policy to decide when and how much to order: when $IP_r \le R_r$, a retailer orders a sufficient multiple of a batch to raise its inventory position above R_r. If an order is placed in period t, define $O_r = R_r - IP_r$, and call the random variable O_r the *overshoot*. Define $\bar{o}_r$ as the maximum overshoot; since $\min\{IP_r^-\} = R_r + 1$, $\bar{o}_r = \bar{d} - 1$.

Since all retailers order in integer multiples of a batch, warehouse demand equals a multiple of a batch too. Therefore, all warehouse variables are measured in batches. ($I_w = 1$ means the warehouse has one batch of inventory.)

The warehouse uses an (R_w, nQ_w) policy to choose its orders: when its inventory position is R_w or lower, it orders a sufficient multiple of Q_w *batches* to raise its inventory position above R_w. Therefore, the warehouse's orders are integer multiples of $Q_r Q_w$ units. Each set of Q_w batches the warehouse orders is called a *system batch* (the system's minimum order quantity is Q_w batches).

Each period the warehouse randomly shuffles the retailer orders and then fills the orders in this sequence. (Of course, orders from previous periods are always filled before orders from the current period.) The shuffling is independent of the retailer's identity and order quantity. This policy is called *random allocation*. The allocation policy matters because it influences the amount of time a retailer expects to wait to receive an ordered batch, which in turn influences the performance metrics (e.g., average inventory, fill rate, etc.) Random allocation treats each retailer equally and it does not require centralized information, i.e., the warehouse does not need to know retailer inventory information. Cachon and Fisher (2000) and Liljenberg (1996) demonstrate that allocation policies which explicitly use centralized information can improve system performance, but they are also more complex. See Graves (1996) for another allocation policy which requires centralized information.

The reorder points R_r and R_w are the decision variables. All other variables and parameters are exogenous. Per period per unit charges at each location can include holding costs h_r and h_w, and a backorder penalty cost at the retailer level p. Fill rates, F_r and F_w, can also be considered in the choice of reorder points.

A summary of the notation is in Appendix A. Unless otherwise noted, capital arabic letters denote either decision variables (e.g., R_r), parameters (e.g., Q_r) or random variables (e.g., O_r). Lower case arabic letters denote the realizations of random variables, e.g., o is a realization of O_r. Subscripts on random variables denote location as well as conditional variables.

2. MODEL EVALUATION

The first step in the exact analysis of the model is the evaluation of the lead time distribution for any batch ordered by any retailer, where the lead time is the number of periods between when a batch is ordered and when it is received. With this distribution it is possible to evaluate each retailer's lead time demand distribution. (In some cases a retailer's lead time demand is not independent of the lead time.) $E[I_r]$ and $E[F_r]$ are evaluated with these distributions. Standard results use $E[I_r]$ to evaluate $E[B_r]$ and $E[I_w]$. From the retailer's lead time distribution it is straightforward to evaluate $E[F_w]$. Finally, the retailer's expected safety stock is evaluated exactly, but an approximation is used to evaluate the warehouse's safety stock.

2.1. Lead Time Analysis

An example helps illustrate the intuition involved in the evaluation of the lead time distribution. In this example there are three retailers, a system batch contains four batches ($Q_w = 4$), the warehouse orders a sufficient number of batches each period to raise its inventory position above three ($R_w = 3$), and the warehouse receives its orders in two periods ($L_w = 2$). Figure 1 displays sample data from periods one through four. Retailer three orders two batches in period one, and the other retailers order one batch each. A batch is represented by a shaded rectangle and a retailer order, which is a set of batches, is represented by an unshaded (and larger) rectangle. Recall that the warehouse randomly sorts the retailer orders each period. In this example, retailer two is served first in period one, followed by retailer three and finally retailer one. To help with their identification, batches are numbered in the sequence the warehouse fills them: batches one through four are ordered in period one, batches five through seven are ordered in period two, etc.

Ovals represent system batches and crossed marked rectangles represent the physical batches. A physical batch fills the ordered batch displayed below it. Recall that a batch is not necessarily shipped in the same period it is ordered. For example, batch ten is filled by the third physical batch in system batch three. If that physical batch is at the warehouse at the start of period four, batch ten is filled in period four, otherwise batch ten is filled in the period after it arrives at the warehouse. When it is shipped depends on when the warehouse orders system batch three. If the system batch is ordered in period one or earlier, it will arrive

Figure 1. Sample sequence of events.

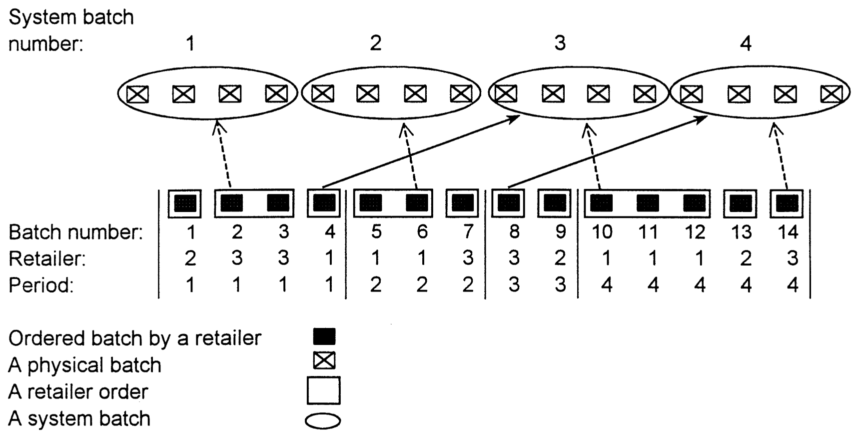

at the warehouse by period three (since $L_w = 2$), and will therefore be available for immediate shipment. However, if the warehouse orders system batch three in period two, those batches will arrive at the end of period four and can only be shipped to retailers in period five (i.e., one period late for batch ten).

When does/did the warehouse order system batch three? The answer depends on the warehouse's reorder point and the realization of retailer orders (which depends on the realization of demands, as is shown later). Suppose at the beginning of period one the warehouse's inventory position is $R_w + Q_w = 7$. After registering batch one, the warehouse's inventory position is six and after registering batch four the warehouse's inventory position is $R_w = 3$. So in period one the warehouse will order one system batch. A *trigger* is an ordered batch that causes the warehouse to order a system batch, e.g., batch four is a trigger.

When the warehouse's inventory position is R_w, there are R_w batches at the warehouse or on route to the warehouse which have not been committed to any ordered batch. Therefore, the vth batch in a system batch fills the $(R_w + v)$th subsequent ordered batch after its trigger. In other words, if batch b is filled by the vth batch in a system batch, then batch $b - (R_w + v)$ is this system batch's trigger. In Figure 1, for the case of $R_w = 3$ system batches and their associated trigger batches are connected with the solid arrows.

When $R_w \geq -1$, a system batch's trigger is never ordered after the first batch in the system batch, i.e., when $R_w \geq -1$, $b - (R_w + 1) \leq b$. In this case, the maximum shipping delay is $L_w + 1$ periods: if a batch is ordered in period t and the trigger batch is ordered in period t too, then the system batch arrives at the warehouse at the end of period $t + L_w$ and batches can be shipped in period $t + L_w + 1$. However, shipping delays can be longer than $L_w + 1$ periods when $R_w < -1$. For example, instead of $R_w = 3$, suppose the warehouse implements $R_w = -3$. If batch b is filled by the first batch in a system batch, then batch $b + 2$ is the trigger, which is the second batch ordered *after* batch b. In Figure 1, given $R_w = -3$, dotted arrows associate system batches with their triggers. For instance, the first batch in system batch three fills batch eight, but the warehouse orders system batch three in period four. In this case batch eight is shipped in period seven, a four period delay. So the shipping delay can be very long when $R_w < -1$.

The intuition in the example is now formalized. Let $\beta_r(o)$ be the number of batches a retailer orders in a period in which it experiences an overshoot o,

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

Let o be retailer i's period t overshoot. (It is assumed that retailer i submits an order in period t.) Let batch b be the jth batch in retailer i's order ($1 \leq j \leq \beta_r(o)$), and suppose batch b is filled by the Vth batch in some system batch, $V \in [1, Q_w]$. V is a random variable, and let v be its realization. Let U_{ojv} be a random variable equal to the number of periods the warehouse delays shipping batch b conditioned on the realizations of O_r and V. In the example, batch eleven ($b = 11$) is the second batch ($j = 2$) in retailer one's ($i = 1$) period four order ($t = 4$) and it is filled by the fourth batch in system batch three ($v = 4$). Also in the example, $U_{ojv} = 0$ when $R_w = 3$, and $U_{ojv} = 3$ when $R_w = -3$.

Batch $b - j + 1$ is the first batch in retailer i's order. So if the trigger is ordered before batch $b - j + 1$, $U_{ojv} \leq L_w + 1$ (because then the trigger is ordered no later than period t). Since batch $b - (R_w + v)$ is the trigger, $U_{ojv} \leq L_w + 1$ when $b - (R_w + v) < b - j + 1$, or

$$0 \leq R_w + v - j. \tag{1}$$

If the above condition does not hold, the trigger is ordered no earlier than the first batch in retailer i's order, so then $U_{ojv} \geq L_w + 1$.

For now, assume that condition (1) holds. Let XB_o^τ equal the number of batches ordered over periods $[t - \tau, t]$, including only batches in period t ordered *before* retailer i's order. Given this definition, batch $b - (XB_o^\tau + j)$ is the last batch ordered in period $t - (\tau + 1)$, $\tau \geq 0$. For $U_{ojv} \leq u$ it must be that the trigger is ordered in period $t - (L_w - u + 1)$ or earlier. This occurs if the last batch ordered in period

$t-(L_w-u+1)$ is greater than or equal to the trigger. Batch $b-(XB_o^{L_w-u}+j)$ is the last batch ordered in period $t-(L_w-u+1)$ and batch $b-(R_w+v)$ is the trigger, so this occurs when

$$b-(XB_o^{L_w-u}+j)\geq b-(R_w+v),$$

which simplifies to $XB_o^{L_w-u}\leq R_w+v-j$. Hence,

$$\Pr(U_{ojv}\leq u|R_w+v-j\geq 0) = \begin{cases}\Pr(XB_o^{L_w-u}\leq R_w+v-j) & 0\leq u\leq L_w,\\ 1 & u\geq L_w+1.\end{cases} \quad (2)$$

Now assume (1) does not hold, so $U_{ojv}\geq L_w+1$. Nevertheless, the evaluation of the shipping delay proceeds in a similar manner. Let XF_o^τ equal the number of batches ordered over periods $[t, t+\tau]$, including only batches in period t ordered *after* retailer i's order. Batch $b-j+\beta_r(o)$ is the last batch in retailer i's order, so batch $b-j+\beta_r(o)+XF_o^\tau+1$ is the first batch ordered in period $t+\tau+1$, $\tau\geq 0$. The trigger is ordered before period $t+\tau+1$ if the first batch ordered in period $t+\tau+1$ is greater than the trigger, i.e.,

$$b-j+\beta_r(o)+XF_o^\tau+1>b-(R_w+v),$$

which simplifies to

$$XF_o^\tau>-1-\beta_r(o)-(R_w+v-j).$$

If the trigger is ordered before period $t+\tau+1$, the shipping delay $L_w+\tau+1$ or fewer periods,

$$\Pr(U_{ojv}\leq u|R_w+v-j<0) = \Pr(XF_o^{u-L_w-1}>-1-\beta_r(o)-(R_w+v-j)). \quad (3)$$

Now define U_{oj} as the shipping delay of retailer's i's jth batch. To evaluate U_{oj} from U_{ojv} requires the distribution for the random variable V as well as an understanding of the relationship between V and O_r (e.g., are they correlated or independent?). The next theorem provides an important result which is used in the proof of the immediate goal, Theorem 2, as well as in several other subsequent results.

THEOREM 1. *IP_r^- is uniformly distributed on the interval $[R_r+1, R_r+Q_r]$, IP_w^- is uniformly distributed on the interval $[R_w+1, R_w+Q_w]$, and the beginning of period inventory positions of the retailers are independent from each other and independent of IP_w^-. (See Appendix* B *for proof.)*

THEOREM 2. *For any retailer order triggered by an overshoot O_r, the jth batch in that order is filled by the Vth batch in some system batch, where V is uniformly distributed on the interval $[1, Q_w]$ and independent of O_r. (See Appendix* B *for proof.)*

It follows from Theorem 2 that

$$\Pr(U_{oj}\leq u)=\frac{1}{Q_w}\sum_{v=1}^{Q_w}\Pr(U_{ojv}\leq u). \quad (4)$$

The average shipping delay, $E[U]$, is evaluated by taking a weighted average over the expected delays of all possible overshoots and batches,

$$E[U]=E_{oj}[E[U_{oj}]] = \frac{\sum_{o=0}^{\bar{o}_r}\sum_{j=1}^{\beta_r(o)}E[U_{oj}]\Pr(O_r=o)}{\sum_{o=0}^{\bar{o}_r}\Pr(O_r=o)\beta_r(o)}. \quad (5)$$

See Appendix C for the evaluation of O_r's distribution function.

It is worthwhile to mention that U_{oj} is not conditioned on the lead time of the other batches in retailer i's order. Therefore, it is not possible to use U_{oj} to evaluate the probability that the second batch in retailer i's order is delayed u or fewer periods given that the first batch is delayed u' periods. Indeed, that evaluation is quite complex because the two batches might be filled from different system batches.

2.2. Retailer Ordering Processes

In §2.1 it is shown how the retailers' ordering processes XF_o^τ and XB_o^τ are used to evaluate the lead time distribution U_{oj}. This section provides a procedure to evaluate these ordering processes.

Recall that retailer i experiences an overshoot o in period t. For retailer i, let IP_o^- be its inventory position at the start of period t and let IP_o^+ be its inventory position at the end of period t (which is also its inventory position at the start of period $t+1$), both conditional on its period t overshoot.

To link demands to orders, define a *trigger demand* as a demand that causes a retailer to order a batch, i.e., when a trigger demand occurs, the retailer's inventory position falls from R_r+1 to R_r. Given that the retailers implement (R_r, nQ_r) policies, each retailer will order a batch every Q_r units of demand, so trigger demands occur every Q_r units of demand. In other words, counting batches and demands after a trigger demand's batch, the bth batch is ordered when the bQ_rth subsequent demand occurs. Suppose a retailer *begins* an interval of periods with an inventory position k and experiences d demands during that interval. Let $Y(k, d)$ be the number of trigger demands that occur in the interval, and hence $Y(k, d)$ is the number of batches the retailer orders in the interval. To develop the $Y(k, d)$ function, identify the last trigger demand that occurs *before* the interval of periods. After that demand and before the interval begins R_r+Q_r-k demands occurred. (After a trigger demand a retailer's inventory position is R_r+Q_r, so its inventory position after the next demand is R_r+Q_r-1, etc.) Hence,

$$Y(k, d)=\left\lfloor\frac{R_r+Q_r-k+d}{Q_r}\right\rfloor.$$

Suppose a retailer *ends* an interval of periods with an inventory position k' but, as before, there are d demands during the interval. There are $k'-R_r-1$ demands which occur before the first trigger demand to occur after the τ periods. Looking backwards in time, after the last trigger demand to occur before the τ periods (which in standard forward looking time is the first trigger demand to occur after the τ periods) through the end of the τ periods, there are $d+k'-R_r-1$ demands, and

$$\left\lfloor \frac{k'-R_r-1+d}{Q_r} \right\rfloor$$

batches ordered during the τ periods. Hence, there are

$$Y(2R_r+1+Q_r-k', d) \tag{6}$$

batches ordered during the τ periods.

Define Y_m^τ as the number of batches m retailers order over an interval of τ periods when they begin these periods with an inventory position in steady state. For one retailer

$$Y_1^\tau = Y(IP_r^-, D^\tau),$$

where recall from Theorem 1 that IP_r^- is uniformly distributed on the interval $[R_r+1, R_r+Q_r]$. Since the retailer's inventory positions are independent, Y_1^τ and Y_{m-1}^τ are independent, and

$$Y_m^\tau = Y_1^\tau + Y_{m-1}^\tau$$

is a simple convolution. The following convenient result implies that the ordering process of retailers in steady state looking forward in time is equivalent to their ordering process looking backwards in time.

THEOREM 3. *Y_m^τ is the number of batches m retailers order in steady state over an interval of τ periods. (See Appendix* B *for proof.)*

Now consider retailer i's ordering processes. Define YF_o^τ as the number of batches retailer i orders over periods $[t+1, t+\tau]$,

$$YF_o^\tau = Y(IP_o^+, D^\tau).$$

Given the period t overshoot o,

$$IP_o^+ = R_r + Q_r - \left(o - \left\lfloor \frac{o}{Q_r} \right\rfloor Q_r\right). \tag{7}$$

Note that IP_o^+ is not a random variable (but YF_o^τ is, because D^τ is a random variable).

Consider retailer i's ordering process before period t. Define YB_o^τ as the number of batches retailer i orders over periods $[t-\tau, t-1]$. From (6),

$$YB_o^\tau = Y(2R_r+1+Q_r-IP_o^-, D^\tau).$$

IP_o^- is not uniformly distributed because it is conditioned on retailer i's period t overshoot. See Appendix C for its evaluation.

Now focus attention on XF_o^τ, which (recall) is the number of batches the retailers order over periods $[t, t+\tau]$, including only batches ordered after retailer i's period t order. Let retailer i be the Mth retailer the warehouse processes in period t. Given that the warehouse fills retailer orders in a random sequence which is independent of the retailer identities and order quantities, $\Pr(M=m) = 1/N$. Define XN^τ as all of the batches in XF_o^τ except for retailer i's batches. Y_{m-1}^τ are the batches ordered over periods $[t+1, t+\tau]$ by the $m-1$ retailers who are processed before retailer i's period t order. $Y_{N-m}^{\tau+1}$ are the remaining batches in XN^τ, i.e., the batches ordered over periods $[t, t+\tau]$ by the $N-m$ retailers who are processed after retailer i's period t order. Hence,

$$\Pr(XN^\tau \le b) = \frac{1}{N}\sum_{m=1}^{N} \Pr(Y_{m-1}^\tau + Y_{N-m}^{\tau+1} \le b). \tag{8}$$

It remains to include retailer i's batches,

$$XF_o^\tau = XN^\tau + YF_o^\tau.$$

From Theorem 3, XN^τ is also all of the batches in XB_o^τ except for retailer i's batches, hence

$$XB_o^\tau = XN^\tau + YB_o^\tau.$$

Note that the mechanics of the warehouse's allocation rule only play a role in the evaluation of XN^τ. Furthermore, the evaluation of XN^τ demonstrates why it is computationally convenient to assume retailers have identical demand processes and minimum order quantities. In (8) it is sufficient to know in period t that $m-1$ retailers are processed before retailer i and $N-m$ retailers are processes after retailer i because the ordering processes of the retailers are identical. However, if there were heterogeneity in the retailers' ordering processes (due to either different minimum order quantities or demand processes), then it would also be necessary to know *which* $m-1$ retailers are processed before retailer i. In that case, far more than N convolutions would be needed to evaluate XN^τ.

2.3. Retailer Lead Time Demand

The previous two sections outline the evaluation of U_{oj}. This distribution is the basis for the evaluation of performance measures of interest. However, before evaluating those performance measures, this section focuses on the relationship between a retailer's lead time distribution and its demand.

Define D_{oju}^τ as retailer i's demand over periods $[t+1, t+\tau]$ when the warehouse delays shipping the jth batch in retailer i's period t order by u periods. When $u \le L_w+1$, the trigger occurs in period t or earlier, so the arrival of retailer i's jth batch is independent of demand after period t, i.e., $D_{oju}^\tau = D^\tau$. This is always true when $R_w \ge -1$ because with those reorder points it is always true (recall) that $u \le L_w+1$.

When $u > L_w+1$, the trigger occurs in period $t+1$ or later. This means that the timing of the trigger is not independent of demand in period $t+1$ or later. To illustrate, consider the example in Figure 1 when $R_0 = -3$. Batch eight is filled with the first batch in system batch three, which is triggered when batch ten occurs. In this example, batch ten happens to occur in period four. But suppose retailers one and two do not order batches in period four. In that case, if batch ten occurs in period four, retailer three must have ordered at least one batch. This means that retailer three's period four demand could not have been zero; if it were zero, retailer three would not have ordered a batch and batch ten would occur only after period four. In short, the higher retailer i's demand in periods $[t+1, t+\tau]$, the more batches retailer i will order in this interval and the more likely the jth batch will arrive sooner (because the trigger is more likely to occur in that interval).

The rest of this section assumes $u > L_w + 1$, so the trigger occurs after period t, i.e., in period $t+u-L_w-1$. Define B_u as the number of batches retailer i orders over periods $[t+1, t+u-L_w-1]$, and b_u as its realization. From (3),

$$\Pr(U_{ojv} \le u | b_u) = \Pr(XF_o^{u-L_w-1} - b_u + b_u > -1 - \beta_r(o) - (R_w + v - j)),$$

but since $XF_o^{u-L_w-1} - B_u = XN^{u-L_w-1}$, the above can be written as

$$\Pr(U_{ojv} \le u | b_u) = \Pr(XN^{u-L_w-1} + b_u > -1 - \beta_r(o) - (R_w + v - j)).$$

It follows that

$$\Pr(U_{ojv} = u | b_u, b_{u-1}) = \Pr(U_{ojv} \le u | b_u) - \Pr(U_{ojv} \le u-1 | b_{u-1}).$$

Like (4),

$$\Pr(U_{oj} = u | b_u, b_{u-1}) = \frac{1}{Q_w} \sum_{v=1}^{Q_w} (U_{ojv} = u | b_u, b_{u-1}).$$

If retailer i has d_u demand over periods $[t+1, t+u-L_w-1]$, it will order $\beta_r(o, d_u)$ batches, where

$$\beta_r(o, d) = \left\lfloor \frac{o+d}{Q_r} \right\rfloor - \left\lfloor \frac{o}{Q_r} \right\rfloor.$$

Since consumer demands are independent across periods, the probability of realizing demands d_{u-1} and d_u is

$$\Pr(D^1 = d_u - d_{u-1}) \Pr(D^{u-L_w-2} = d_{u-1}).$$

Therefore, from Bayes' theorem,

$$\Pr\left(D_{oju}^{u-L_w-1} = d_u\right) = \left(\sum_{d_{u-1}=0}^{d_u} \Pr(D^1 = d_u - d_{u-1}) \Pr(D^{u-L_w-2} = d_{u-1}) \cdot \Pr(U_{oj} = u | \beta_r(o, d_u), \beta_r(o, d_{u-1})) \right) \Big/ (\Pr(U_{oj} = u)).$$

Since retailer i's consumer demand after period $t+u-L_w-1$ is independent of u,

$$D_{oju}^{\tau} = \begin{cases} D_{oju}^{u-L_w-1} + D^{\tau-(u-L_w-1)} & \tau \ge u - L_w - 1;\ u > L_w + 1, \\ D^{\tau} & u \le L_w + 1. \end{cases}$$

2.4. Retailer Average Inventory

Let $E[S]$ be the expected number of periods a unit is recorded in inventory (a unit's expected sojourn in inventory). From Little's Law, $E[I_r] = \mu_r E[S]$, where $\mu_r = E[D^1]$.

What is the sojourn of the cth unit in retailer i's jth batch from its period t order? To answer this question, consider a simpler scenario. As defined in §2.2, a trigger demand is one that causes a retailer to order a batch. Consider the trigger demand for the first batch in retailer i's period t order. This demand lowers retailer i's inventory position from $R_r + 1$ to R_r, which means that if there were no additional batches ordered, the retailer would only be able to fill R_r more demands. Hence, the first unit in the first batch retailer i orders in period t must fill the (R_r+1)th demand relative to the trigger demand. In general, the cth unit in the jth batch must fill the $(R_r + c + (j-1)Q_r)$th demand relative to the trigger demand. The jth batch arrives at retailer i in period $t+u+L_r$, where u is the warehouse's shipping delay. So the cth unit in retailer i's jth batch from its period t order is recorded in inventory for $s > 0$ periods when the $(R_r + c + (j-1)Q_r)$th demand relative to the trigger demand occurs in period $t+u+L_r+1+s$. (A unit arriving in period $t+u+L_r$ is not recorded in that period's inventory because inventory is measured within a period after demand, but before replenishments arrive. Further, if the unit is sold in period $t+u+L_r+1$, it will also not be recorded in inventory for that period.)

By definition of the overshoot, retailer i's inventory position is $R_r - o$ after demand in period t. Hence, in period t there were o additional demands after the trigger demand for retailer i's first batch. Therefore, the $(R_r + c + (j-1)Q_r)$th demand relative to the trigger demand occurs in period $t+u+L_r+1+s$ or earlier if at least $R_r + c + (j-1)Q_r - o$ demands occur over periods $[t+1, t+u+L_r+1+s]$, i.e.,

$$\Pr(S_{ojuc} \le s) \Pr\left(D_{oju}^{u+L_r+1+s} \ge R_r + c + (j-1)Q_r - o\right).$$

It follows that

$$E[S_{ojuc}] = \sum_{s=0}^{\infty} 1 - \Pr(S_{ojuc} \le s) = \sum_{s=0}^{\infty} \Pr\left(D_{oju}^{u+L_r+1+s} < R_r + c + (j-1)Q_r - o\right).$$

See Appendix D for details on how to evaluate $E[S_{ojuc}]$ with finite effort. Unconditioning yields

$$E[S] = E_{ojuc}[E[S_{ojuc}]] = \left(\frac{1}{Q_r} \sum_{o=0}^{\bar{o}_r} \sum_{j=1}^{\beta_r(o)} \sum_{u=0}^{\bar{u}} \sum_{c=1}^{Q_r} E[S_{ojuc}] \Pr(O_r = o) \Pr(U_{oj} = u) \right) \Big/ \left(\sum_{o=0}^{\bar{o}_r} \Pr(O_r = o) \beta_r(o) \right), \tag{9}$$

where $\bar{u}$ is the maximum shipping delay.

2.5. Retailer Backorder Level

From the definition of a retailer's inventory position,

$$E[IP_r] = E[I_r] - E[B_r] + E[IT_r].$$

Since a retailer's inventory position is uniform on the interval $[R_r + 1, R_r + Q_r]$ in steady state (Theorem 1),

$$E[IP_r] = R_r + \frac{Q_r + 1}{2}.$$

From Little's Law, $E[IT_r] = \mu_r(E[U] + L_r + 1)$. Hence,

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

2.6. Retailer Fill Rate

Let F_{ojuc} be the probability the retailer fills immediately from stock the cth unit in the jth batch of retailer i's period t order, given that o is this order's overshoot and the warehouse delays shipping the jth batch by u periods. Following the procedure in §2.4, this unit fills the $(R_r + c + (j-1)Q_r)$th demand after the trigger demand for the first batch in retailer i's period t order. This demand occurs in period $t + u + L_r$ or earlier if at least $R_r + c + (j-1)Q_r - o$ demands occur over periods $[t+1, t+u+L_r]$, and if this occurs, then the cth unit arrives too late to fill its demand (because it arrives at the end of period $t + u + L_r$). Hence,

$$F_{ojuc} = 1 - \Pr\left(D_{oju}^{u+L_r} \geq R_r + c + (j-1)Q_r - o\right).$$

Overall, the retailers' expected fill rate is $E[F_r] = E_{ojuc} \cdot [E[F_{ojuc}]]$, an evaluation which is analogous to (9).

2.7. Retailer Safety Stock

A retailer's safety stock is often defined as its average net inventory (on hand inventory minus backorders) just before a replenishment arrives. However, in this setting this definition is ambiguous. For example, suppose a retailer orders two batches but the batches arrive in different periods. Should safety stock be measured in each of the periods these batches arrive, or should it be measured only in the period the first batch arrives? The procedure presented in this research cannot evaluate the former because that requires evaluating the arrival time of the second batch conditional on the first. Hence, the former definition is adopted. (The two definitions are the same when a retailer orders only one batch per order, which becomes more likely as Q_r increases.)

Again suppose retailer i places an order in period t due to an overshoot o. Let A_{ou} be the safety stock of this order when the warehouse delays shipping the first batch u periods. Before the order is recorded, the retailer's inventory position is $R_r - o$. Since the first batch in the order arrives in period $t + u + L_r$, the retailer's net inventory before this arrival will be $R_r - o - D_{oju}^{u+L_r}$, so

$$A_{ou} = R_r - o - D_{oju}^{u+L_r},$$

where $j = 1$. Averaging over possible overshoots and shipping delays,

$$E[A] = \sum_{o=0}^{\bar{o}_r} \sum_{u=0}^{\bar{u}} E[A_{ou}] \Pr(O_r = o) \Pr(U_{oj} = u),$$

where $j = 1$.

2.8. Warehouse Average Backorder and Inventory

The average number of batches a retailer has ordered but the warehouse has not shipped is $\mu_r E[U]/Q_r$. Summing across the N retailers yields the warehouse's average backorder, measured in batches,

$$E[B_w] = \frac{N}{Q_r} \mu_r E[U] = \mu_w E[U],$$

where $\mu_w = N\mu_r / Q_r$.

From the definition of a firm's inventory position,

$$E[IP_w] = E[I_w] - E[B_w] + E[IT_w].$$

From Theorem 1, IP_w is uniformly distributed on the interval $[R_w + 1, R_w + Q_w]$, so

$$E[IP_w] = R_w + \frac{Q_w + 1}{2}.$$

The warehouse's source always delivers inventory in L_w periods, so from Little's Law, $E[IT_w] = \mu_w(L_w + 1)$. (A batch ordered after demand in period t, but before inventory is measured, is delivered after inventory is measured in period $t + L_w$, so it is recorded as on order for periods $[t, t + L_w]$.) Combining the above results,

$$E[I_w] = R_w + \frac{Q_w + 1}{2} + E[B_w] - \mu_w(L_w + 1).$$

2.9. Warehouse Fill Rate

Let F_{oj} equal the probability that the warehouse does not delay the shipment of retailer i's jth batch in period t, i.e.,

$$F_{oj} = \Pr(U_{oj} = 0).$$

Averaging over all overshoots and batches yields the warehouse's average fill rate, $E[F_w] = E_{oj}[F_{oj}]$, an evaluation which is analogous to (5).

2.10. Warehouse Safety Stock and Cycle Stock-Out Probability

The warehouse's safety stock, A_w, is the net inventory (on hand inventory minus backorders) at the warehouse when it receives a replenishment. Let O_w be the warehouse's overshoot when it places an order. Suppose an order is placed in period t due to an overshoot $O_w = o$. This order will arrive in period $t + L_w$. Let $\widetilde{Y}_o$ be the number of batches ordered by the retailers over periods $[t+1, t+L_w]$, conditional on the warehouse's overshoot in period t, o. The warehouse's safety stock is then

$$A_w = R_w - o - \widetilde{Y}_o$$

and its expected safety stock is

$$E[A_w] = R_w - \sum_{o=0}^{\bar{o}_w} \Pr(O_w = o)(o + E[\widetilde{Y}_o])$$
$$= R_w - E[O_w] - E_o[E[\widetilde{Y}_o]],$$

where $\bar{o}_w$ is the warehouse's maximum overshoot. See Appendix C for the evaluation of O_w.

The exact evaluation of $E[\widetilde{Y}_o]$ can be quite cumbersome. Therefore, the following approximation is proposed:

$$E[A] \simeq R_w - E[O_w] - \mu_w L_w.$$

This approximation assumes that the retailer's ordering process over periods $[t+1, t+L_w]$ is independent of the ordering process in period t. The above is exact when $Q_r = 1$, because in that situation it is indeed true that a retailer's future orders are independent of its current orders. (When $Q_r = 1$ a retailer's order equals its demand, and by assumption its demands are independent across periods.)

Define the warehouse's cycle stock-out probability as the probability the warehouse will have some backorder between the time it places an order and the time it receives an order. This probability is

$$\sum_{o=0}^{\bar{o}_w} \Pr(O_w = o)\Pr(\widetilde{Y}_o > R_w - o).$$

Since, as described above, $\widetilde{Y}_o$ is difficult to evaluate, the following approximation is proposed:

$$\sum_{o=0}^{\bar{o}_w} \Pr(O_w = o)\Pr\big(Y_N^{L_w} > R_w - o\big), \tag{10}$$

where, recall that $Y_N^{L_w}$ is the number of batches N retailers order over L_w periods when they begin these periods with their inventory position in steady state. $Y_N^{L_w}$ underestimates the warehouse's demand over periods $[t+1, t+L_w]$ when O_w draws a small realization, because low demand in period t predicts higher demand over later periods, whereas $Y_N^{L_w}$ overestimates demand when O_w draws a large realization. However, when $Q_r = 1$, (10) is indeed exact because, as mentioned above, in that case the retailer's ordering process is independent across periods.

3. SELECTING POLICIES

Define C_r and C_w as the holding and backorder costs charged to the retailers and the warehouse respectively:

$$C_r = h_r N E[I_r] + pNE[B_r];$$
$$C_w = h_w Q_r E[I_w].$$

If firms coordinated the selection of the reorder points (or if a single firm controlled the entire supply chain), then policies could be chosen to minimize system wide costs:

$$\min_{R_r, R_w} C_r + C_w. \tag{11}$$

For a given R_w, it is simple to show that C_r is convex in R_r, so the solution to (11) for a given R_w is simple. Unfortunately, $C_r + C_w$ is not (necessarily) jointly convex in R_r and R_w. Hence, it is necessary to search for the cost minimizing R_w. To limit the search range, note that $R_w < -Q_w$ is never optimal because $E[I_w | R_w \le -Q_w] = 0$. Furthermore,

$$R_w \ge N \left\lfloor \frac{\bar{d}(L_w+1) + Q_r - 1}{Q_r} \right\rfloor$$

is never optimal because then the warehouse never delays shipping a batch. (Any increase in R_w would then raise warehouse inventory but have no effect on $E[I_r]$ or $E[B_r]$.)

Policies could also be selected to minimize system wide inventory subject to a constraint that the average retailer fill rate equals α or higher:

$$\min_{R_r, R_w} h_r E[I_r] + h_w Q_r E[I_w]$$
$$\text{s.t} \quad E[F_r] \ge \alpha. \tag{12}$$

The procedure to solve (11) applies in the solution of (12): search over the plausible range for R_w, and for each R_w find the inventory minimizing R_r.

Computational effort depends on several factors. The retailer ordering processes (XN^τ, Y_m^τ, etc.) need only be evaluated once per model since they are independent of the reorder points. Among those processes, the most computationally intensive is XN^τ which is $O(\tau^3 \bar{b}^2 N^3)$, where $\bar{b}$ is the most number of batches a single retailer will order in a single period,

$$\bar{b} = \left\lfloor \frac{Q_r + 1 + \bar{d}}{Q_r} \right\rfloor.$$

(Note that O is not referring to the overshoot, but rather to the computational effort of the evaluation.) For how large a τ is it necessary to evaluate XN^τ? From (2) and (3), XN^τ needs to be evaluated for up to $\tau = \max\{L_w, \bar{u} - L_w - 1\}$. Recall that $\bar{u} = L_w + 1$ when $R_w \ge -1$, so in that case up to XN^{L_w} needs to be evaluated. However, when $R_w = -Q_w$, and the retailer order process is slow relative to Q_w (i.e., many periods are required before the retailers order Q_w batches), then $\bar{u}$ can be significantly larger than $2L_w + 1$, thereby increasing the computational effort. Hence, there is not a clear relationship between computational effort and Q_r or N. It holds that $\bar{b}$ is decreasing in Q_r, thereby reducing computational effort, but $\bar{u}$ is increasing in Q_r. Similarly, holding $\bar{u}$ constant, effort to evaluate XN^τ is increasing in N, but in fact, $\bar{u}$ is decreasing in N, so the net effect is unclear.

4. NUMERIC EXAMPLES

Eighty scenarios are presented to demonstrate the utility of this model. The first 32 scenarios assume Poisson consumer demand at each retailer and each combination from the following sets of parameters: $\mu_r \in \{0.1, 1\}$; $N \in \{4, 32\}$; $Q_r \in \{1, 4\}$; $Q_w \in \{1, 4\}$; $L_r = 1$; $L_w = 1$; $h_r = 1$; $h_w = 1$; and $p \in \{5, 20\}$. Svoronos and Zipkin (1988) and Axsäter (1993) also studied these problems, except they assumed the firms operate in an environment with continuous review. This study assumes periodic review with a period length equal to one.

Scenarios 33–48 replicate scenarios 17–32, except the warehouse's lead time equals five periods, $L_w = 5$. Scenarios 49–64 replicate problems 17–32, except consumer demand in these problems follows a discrete version of the normal distribution:

$$\Pr(D^1 = d) = \begin{cases} \Phi(0.5) & d = 0, \\ \Phi(d+0.5) - \Phi(d-0.5) & d > 0, \end{cases}$$

where $\Phi(\cdot)$ is the normal distribution function with $\mu_r = 1$ and a standard deviation equal to 0.5. The last set of 16 scenarios also replicate problems 17–32, except consumer demand in these problems follows a negative binomial distribution,

$$\Pr(D^1 = d) = \binom{d+r-1}{d} q^r (1-q)^d, \qquad d = 0, 1, \ldots,$$

with parameters $q = 0.5$ and $r = 1$. With this distribution $\mu_r = 1$ and the variance of consumer demand equals 2. For Poisson demand with $\mu_r = 0.1$, $\bar{d} = 3$, and with $\mu_r = 1$, $\bar{d} = 7$. With normal demand $\bar{d} = 3$ and with negative binomial demand $\bar{d} = 13$. In all of the $Q_w > 1$ scenarios, $\bar{u}$ is chosen such that $\Pr(U_{oj} \leq \bar{u}) \geq 0.99999$ for all o and j.

Note that scenarios 17–32 and 49–80 differ only in the variability of consumer demand: The coefficient of variation equals 0.5 for the normal, equals 1 for the Poisson and equals (approximately) 1.44 for the negative binomial. Table I displays the parameter values for each scenario.

Tables II and III present the reorder policies for the warehouse and retailers that minimize total system costs (inventory holding plus backorder costs) for each of the scenarios. (The author's webpage, ⟨opim.wharton.upenn.edu/∼cachon⟩, contains the compiled code, instructions and sample scenarios.) The tables also list for each location, assuming the optimal policies are selected, average inventory, average backorders, average fill rates and average safety stock. In addition, the warehouse's stock-out probability is listed. These tables reveal that optimal policies do not necessarily guarantee high retailer service levels: retailer fill rates range from 63% to 98%. In addition, there is considerable variation in the optimal warehouse fill rate: values range from 0% to 99%. However, the optimal warehouse safety stock tends to increase in the backorder penalty cost: in 37 of 40 scenarios the optimal warehouse safety stock with $p = 20$ is no lower than the optimal warehouse safety stock with $p = 5$.

As an alternative to minimizing system wide inventory costs, Table IV presents the policies that minimize supply chain total inventory while maintaining at least a 99% retailer fill rate. Again, optimal warehouse fill rates vary considerably (between 0% and 99%).

It would be worthwhile to determine if there is a simple heuristic for choosing the warehouse reorder point which would yield costs reasonably close to optimal (assuming R_r is chosen to minimize costs given R_w). Five heuristics are considered. The first sets $R_w = -Q_w$, which ensures that the warehouse carries no inventory, i.e., the warehouse acts merely as a transit point, or cross docking facility. The second, first suggested by Schwarz, Deuermeyer, and Badinelli (1985) recommends choosing R_w such that warehouse safety stock is approximately $-Q_w$. Graves (1996) found that the warehouse should almost always have nonpositive safety stock, so the third heuristic sets the warehouse safety stock to approximately zero. (Due to the integral constraint on R_w, in heuristics two and three warehouse reorder points are found that yield safety stocks as close as possible to the targeted levels.) The fourth heuristic chooses the best R_w such that the warehouse offers at least a 99% fill rate. This heuristic is guided by practice; it has been observed that practitioners often desire 99% or higher fill rates at all locations in the supply chain. (This observation is based on personal experience with a supplier, a wholesaler and a retailer in the grocery industry.)

The fifth heuristic implements the reorder points that are optimal if the system actually operated under continuous review. The methods to evaluate reorder point policies developed in this paper only apply under periodic review, but there is other research that evaluates periodic review policies with continuous review, e.g., Axsater (1993). With just one exception, the appropriate continuous review model to evaluate has the same parameters as the periodic review model, e.g., the retailers' holding cost per unit per unit time is still h_r, the demand rate per unit of time for each retailer is still μ_r (assuming a period equals one unit of time), etc. The one exception is the warehouse's lead time. That lead time should be chosen so that with either the periodic review or the continuous review model the warehouse has the same average demand between when the warehouse submits an order and when the units in that order begin to incur holding costs. With periodic review a warehouse order placed in period t is received in period $t + L_w$ and first incurs holding costs in period $t + L_w + 1$. So the average demand on the warehouse in that interval is $(L_w + 1)N\mu_r$ units. Since the warehouse's average demand per unit time in the continuous review model is N_{μ_r} units, the appropriate warehouse lead time for the continuous review model is $L_w + 1$ units of time.

Table V presents the increase in costs when using one of the first four heuristics. Generally, the extreme heuristics (one and four) increase supply chain costs substantially, but not always. The third heuristic, setting warehouse safety stock to zero, frequently provides near

Table 1. Scenario parameter values.

Scenario Number	Demand Distribution	μ_r	N	p	L_w	Q_r	Q_w	Scenario Number	Demand Distribution	μ_r	N	p	L_w	Q_r	Q_w
1	Poisson	0.1	4	20	1	1	1	41	Poisson	1	32	20	5	1	1
2	Poisson	0.1	4	20	1	1	4	42	Poisson	1	32	20	5	1	4
3	Poisson	0.1	4	20	1	4	1	43	Poisson	1	32	20	5	4	1
4	Poisson	0.1	4	20	1	4	4	44	Poisson	1	32	20	5	4	4
5	Poisson	0.1	4	5	1	1	1	45	Poisson	1	32	5	5	1	1
6	Poisson	0.1	4	5	1	1	4	46	Poisson	1	32	5	5	1	4
7	Poisson	0.1	4	5	1	4	1	47	Poisson	1	32	5	5	4	1
8	Poisson	0.1	4	5	1	4	4	48	Poisson	1	32	5	5	4	4
9	Poisson	0.1	32	20	1	1	1	49	Normal	1	4	20	1	1	1
10	Poisson	0.1	32	20	1	1	4	50	Normal	1	4	20	1	1	4
11	Poisson	0.1	32	20	1	4	1	51	Normal	1	4	20	1	4	1
12	Poisson	0.1	32	20	1	4	4	52	Normal	1	4	20	1	4	4
13	Poisson	0.1	32	5	1	1	1	53	Normal	1	4	5	1	1	1
14	Poisson	0.1	32	5	1	1	4	54	Normal	1	4	5	1	1	4
15	Poisson	0.1	32	5	1	4	1	55	Normal	1	4	5	1	4	1
16	Poisson	0.1	32	5	1	4	4	56	Normal	1	4	5	1	4	4
17	Poisson	1	4	20	1	1	1	57	Normal	1	32	20	1	1	1
18	Poisson	1	4	20	1	1	4	58	Normal	1	32	20	1	1	4
19	Poisson	1	4	20	1	4	1	59	Normal	1	32	20	1	4	1
20	Poisson	1	4	20	1	4	4	60	Normal	1	32	20	1	4	4
21	Poisson	1	4	5	1	1	1	61	Normal	1	32	5	1	1	1
22	Poisson	1	4	5	1	1	4	62	Normal	1	32	5	1	1	4
23	Poisson	1	4	5	1	4	1	63	Normal	1	32	5	1	4	1
24	Poisson	1	4	5	1	4	4	64	Normal	1	32	5	1	4	4
25	Poisson	1	32	20	1	1	1	65	Neg. Bin.	1	4	20	1	1	1
26	Poisson	1	32	20	1	1	4	66	Neg. Bin.	1	4	20	1	1	4
27	Poisson	1	32	20	1	4	1	67	Neg. Bin.	1	4	20	1	4	1
28	Poisson	1	32	20	1	4	4	68	Neg. Bin.	1	4	20	1	4	4
29	Poisson	1	32	5	1	1	1	69	Neg. Bin.	1	4	5	1	1	1
30	Poisson	1	32	5	1	1	4	70	Neg. Bin.	1	4	5	1	1	4
31	Poisson	1	32	5	1	4	1	71	Neg. Bin.	1	4	5	1	4	1
32	Poisson	1	32	5	1	4	4	72	Neg. Bin.	1	4	5	1	4	4
33	Poisson	1	4	20	5	1	1	73	Neg. Bin.	1	32	20	1	1	1
34	Poisson	1	4	20	5	1	4	74	Neg. Bin.	1	32	20	1	1	4
35	Poisson	1	4	20	5	4	1	75	Neg. Bin.	1	32	20	1	4	1
36	Poisson	1	4	20	5	4	4	76	Neg. Bin.	1	32	20	1	4	4
37	Poisson	1	4	5	5	1	1	77	Neg. Bin.	1	32	5	1	1	1
38	Poisson	1	4	5	5	1	4	78	Neg. Bin.	1	32	5	1	1	4
39	Poisson	1	4	5	5	4	1	79	Neg. Bin.	1	32	5	1	4	1
40	Poisson	1	4	5	5	4	4	80	Neg. Bin.	1	32	5	1	4	4

optimal results, but does perform poorly in a few scenarios, especially when there are only a few retailers and demand is slow. The second heuristic, setting warehouse safety stock to $-Q_w$, produces similar results to heuristic three. Overall, it does not appear that any of these heuristics provides a good solution in all scenarios, but for a significant fraction of scenarios the optimal warehouse safety stock is nonpositive. Table VI presents the results for each of the four heuristics when the supply chain inventories are minimized while maintaining at least a 99% retailer fill rate. In this situation, heuristics two and three do not perform as well, but there are still several scenarios in which they are near optimal. Finally, it should be noted that all of these results assume $h_w = h_r$. Clearly, the optimal warehouse reorder point will tend to increase as h_w declines (improving the performance of the high fill rate heuristic). Indeed, if it were free to hold inventory at the warehouse, the optimal solution would have the warehouse offering a 100% fill rate (and heuristics one–three might perform poorly).

Table VII presents the performance of the optimal continuous review policies when they are applied in the periodic review environment. The methodology in Axsater (1993) is used to find the optimal continuous review policies and the methodology in this paper evaluates expected costs with those policies. For high mean consumer demand ($\mu_r = 1$) the continuous review policies increase costs between 18–82% relative to the best periodic review policies. The comparable increase for the low mean consumer demand ($\mu_r = 0.1$) is 0–36%. Interestingly, in all examples the continuous review retailer reorder point is no greater than the optimal reorder point and often smaller, which suggests that the continuous review model underestimates the amount of inventory needed at the retail level.

Table 2. Policies that minimize total supply chain holding and backorder costs (scenarios 1–40).

Scenario			Expected values*									
			Total	Inventory		Backorders		Safety Stock		Fill Rate		Ware. Stockout
Number	R_w	R_r	Cost	Rets.	Ware.	Rets.	Ware.	Rets.	Ware.	Rets. (%)	Ware. (%)	Prob. (%)
1	0	0	6.23	3.09	0.45	0.13	0.25	−1.05	−0.61	81.1	55.2	45
2	0	0	6.87	3.21	1.78	0.09	0.08	−0.88	−0.60	84.5	85.0	45
3	−1	0	10.08	8.48	0.00	0.08	0.80	−1.40	−4.55	91.3	0.0	100
4	−1	0	15.16	9.02	5.42	0.04	0.21	−0.81	−4.55	94.9	72.2	100
5	−1	0	4.09	2.68	0.00	0.28	0.80	−1.60	−1.61	70.5	0.0	100
6	−2	0	4.57	2.75	0.43	0.28	0.73	−1.53	−2.60	72.4	35.7	100
7	−1	−1	7.28	4.88	0.00	0.48	0.80	−5.40	−4.55	66.3	0.0	100
8	−2	−1	11.65	4.66	2.62	0.87	1.41	−6.02	−8.55	63.1	47.2	100
9	7	0	41.77	25.87	2.03	0.69	0.43	−6.82	1.46	85.0	86.9	20
10	6	0	42.05	25.90	2.48	0.68	0.38	−6.77	1.08	85.1	88.4	24
11	0	0	79.23	70.81	0.77	0.38	3.17	−7.97	−4.96	93.8	31.4	71
12	−1	0	80.82	71.16	2.41	0.36	2.81	−7.60	−8.75	94.1	46.9	100
13	4	0	30.27	24.80	0.41	1.01	1.81	−8.21	−1.54	81.5	48.8	64
14	2	0	30.45	24.46	0.36	1.12	2.26	−8.66	−2.92	80.4	40.5	85
15	0	−1	55.86	41.20	0.77	2.78	3.17	−39.97	−4.96	68.8	31.4	71
16	−1	−1	57.19	41.46	2.41	2.66	2.81	−39.60	−8.75	69.1	46.9	100
17	7	4	16.50	11.10	1.12	0.21	1.12	6.88	−0.07	95.3	72.9	41
18	6	4	16.69	11.22	1.48	0.20	0.98	7.02	−0.25	95.6	76.4	43
19	1	3	20.32	12.69	1.61	0.30	1.61	4.32	−1.87	94.0	65.3	45
20	−1	4	22.39	14.74	1.57	0.30	3.57	6.37	−9.52	94.3	45.3	100
21	6	3	11.28	7.05	0.66	0.71	1.66	2.34	−1.07	85.7	60.5	56
22	5	3	11.48	7.20	0.96	0.66	1.46	2.54	−1.25	86.6	65.5	58
23	1	2	14.22	9.09	1.61	0.70	1.61	0.32	−1.87	86.8	65.3	45
24	−1	2	15.95	7.76	1.57	1.32	3.57	−1.63	−9.52	78.9	45.3	100
25	64	4	118.39	94.30	3.72	1.02	2.72	61.29	1.00	97.1	91.5	42
26	63	4	118.46	94.46	4.04	1.00	2.54	61.47	1.50	97.1	92.1	39
27	15	3	140.09	108.53	4.81	1.34	4.81	42.64	0.00	96.3	85.0	43
28	14	3	140.77	109.08	6.20	1.27	4.20	43.25	1.77	96.5	86.9	39
29	66	2	82.07	37.91	4.93	7.85	1.93	−1.93	3.00	79.0	94.0	32
30	64	2	82.14	37.80	4.64	7.94	2.14	−2.14	2.50	78.8	93.3	35
31	13	2	100.01	74.87	1.80	4.67	9.80	5.65	−8.00	88.4	69.5	66
32	11	2	100.36	73.48	1.58	5.06	11.57	3.87	−10.23	87.6	64.1	71
33	25	4	19.27	11.12	3.13	0.25	1.13	6.87	1.93	94.9	77.1	30
34	24	4	19.43	11.22	3.52	0.23	1.02	6.98	1.75	95.2	79.4	31
35	5	4	22.20	15.97	2.23	0.20	2.23	7.70	−1.87	96.3	61.5	48
36	4	4	24.29	16.20	4.00	0.20	2.00	7.93	−5.52	96.4	68.5	72
37	23	3	13.30	6.94	1.95	0.88	1.95	2.05	−0.07	84.0	63.5	45
38	21	3	13.45	6.74	1.75	0.99	2.25	1.75	−1.25	82.6	59.3	55
39	5	2	15.61	8.70	2.23	0.93	2.23	−0.30	−1.87	84.3	61.5	48
40	3	3	17.36	10.90	1.99	0.89	3.99	1.94	−9.52	86.7	47.7	90

*Retailer values are totals for all retailers, warehouse values are in units, "Ware. Stockout Prob" is the probability a backorder occurs between the period a system batch is ordered and the period it arrives.

5. DISCUSSION

This research provides exact results by evaluating the lead time distribution for each batch a retailer orders. Other researchers have obtained exact results in continuous review models using different approaches. Axsäter (1990) evaluates one-for-one ordering policies in a model with Poisson demand. As in this research, he also evaluates the expected holding and back-order costs for each unit ordered by the warehouse and then averages over all possible units. His technique is linked to the assumption of Poisson demand; as a result he does not need to evaluate the lead time distribution. Axsäter (1993) demonstrates that exact results with batch ordering can be obtained by taking a weighted average of the performance of systems with one-for-one ordering. It does not appear that this is a fruitful approach in a periodic review setting. In Axsäter's continuous review model the overshoot is always zero whether one-for-one ordering or batch ordering is implemented (i.e., orders are always placed when the inventory position is exactly $R+1$), so the needed demand and ordering processes are not conditioned on the overshoot. But in periodic review the needed demand distributions are conditioned on the overshoot and the overshoot is not independent of the batch size. Axsäter (2000) and Chen and Zheng (1997) obtain exact results by evaluating the steady state distribu-

Table 3. Policies that minimize total supply chain holding and backorder costs (scenarios 41–80).

Scenario			Expected values*									
			Total	Inventory		Backorders		Safety Stock		Fill Rate		Ware. Stockout
Number	R_w	R_r	Cost	Rets.	Ware.	Rets.	Ware.	Rets.	Ware.	Rets. (%)	Ware. (%)	Prob. (%)
41	194	4	123.86	93.02	7.17	1.18	4.17	59.84	3.00	96.7	87.0	40
42	192	4	123.90	92.82	6.90	1.21	4.39	59.61	2.50	96.6	86.3	41
43	48	3	144.09	108.59	8.75	1.34	4.75	42.70	4.00	96.3	85.3	37
44	46	3	144.60	107.62	7.84	1.46	5.84	41.61	1.77	96.0	82.1	42
45	185	3	85.69	59.61	3.01	4.62	9.01	23.00	−6.00	88.0	72.2	65
46	183	3	85.73	59.35	2.86	4.70	9.36	22.65	−6.50	87.8	71.1	66
47	45	2	102.29	73.76	3.29	5.05	11.29	4.16	−8.00	87.7	66.0	63
48	43	2	102.60	72.53	2.90	5.43	12.90	2.55	−10.23	86.9	61.7	67
49	6	3	8.56	6.81	0.24	0.08	1.25	5.46	−1.01	98.1	68.8	63
50	5	3	8.83	7.00	0.54	0.06	1.05	5.66	−0.83	98.4	73.9	56
51	−1	4	12.40	10.09	0.00	0.12	8.01	3.34	−9.86	97.3	0.0	100
52	−1	3	16.12	10.68	1.50	0.20	3.51	3.84	−9.51	95.7	46.8	100
53	7	2	5.81	3.67	0.61	0.30	0.62	2.08	−0.01	92.8	84.5	38
54	5	2	6.14	3.38	0.54	0.44	1.05	1.66	−0.83	89.6	73.9	56
55	−1	3	9.19	6.51	0.00	0.54	8.01	−0.66	−9.86	88.3	0.0	100
56	−1	2	11.91	7.14	1.50	0.65	3.51	−0.16	−9.51	87.4	46.8	100
57	68	2	57.40	32.57	5.23	0.98	0.32	21.33	4.91	97.0	99.0	12
58	66	2	57.64	32.50	4.84	1.01	0.43	21.22	4.41	96.9	98.7	15
59	15	2	93.84	76.06	4.44	0.67	4.52	22.28	−0.09	98.0	85.9	44
60	14	2	94.62	76.59	5.84	0.61	3.93	22.87	1.69	98.2	87.7	39
61	63	2	40.65	31.54	1.76	1.47	1.85	19.80	−0.09	95.6	94.2	46
62	61	2	40.77	31.32	1.58	1.57	2.16	19.48	−0.59	95.3	93.2	51
63	15	1	70.15	47.11	4.44	3.72	4.52	−9.72	−0.09	89.5	85.9	44
64	13	1	70.69	46.12	3.85	4.14	5.94	−11.14	−2.31	88.5	81.5	50
65	7	6	26.87	18.76	1.57	0.33	1.57	10.44	−0.27	94.2	65.5	42
66	6	6	27.04	18.90	1.92	0.31	1.42	10.60	−0.75	94.4	68.9	47
67	1	5	28.95	20.45	1.89	0.33	1.88	9.06	−2.09	94.3	61.0	45
68	0	5	30.92	20.57	3.77	0.33	1.76	9.18	−5.77	94.4	66.9	77
69	5	4	17.39	10.56	0.68	1.23	2.67	1.34	−2.27	81.0	45.0	63
70	4	4	17.51	10.74	0.93	1.17	2.42	1.59	−2.75	81.8	50.0	69
71	1	2	19.14	9.64	1.89	1.52	1.88	−2.94	−2.09	76.8	61.0	45
72	−1	3	20.60	11.71	1.73	1.43	3.72	−0.78	−9.77	79.9	42.8	100
73	68	5	192.96	128.35	7.52	2.85	2.51	61.59	5.01	93.1	92.1	30
74	66	5	193.01	128.20	7.20	2.88	2.69	61.41	4.51	93.1	91.6	32
75	16	4	206.18	142.93	7.86	2.77	3.85	51.72	4.00	93.5	88.0	33
76	15	4	206.95	143.32	9.41	2.71	3.40	52.17	5.69	93.6	89.4	30
77	63	3	121.24	69.05	4.51	9.54	4.50	−4.39	0.01	79.1	85.9	46
78	61	3	121.29	68.87	4.27	9.63	4.77	−4.66	−0.49	78.9	85.1	48
79	14	2	133.62	81.85	3.76	9.60	7.76	−16.19	−4.00	79.9	75.9	54
80	13	2	134.05	82.45	4.90	9.34	6.89	−15.32	−2.31	80.3	78.6	49

*Retailer values are totals for all retailers, warehouse values are in units, "Ware. Stockout Prob" is the probability a backorder occurs between the period a system batch is ordered and the period it arrives.

tion of the firms' net inventory (on hand inventory minus backorders). This requires evaluating the steady state distribution of the number of batches the warehouse has backordered for each retailer, which may be computationally quite burdensome in a periodic review environment with general demand distributions. However, if that distribution were known it would not be necessary to evaluate D^{τ}_{oju}, thereby saving some computational effort.

The model in this paper is general enough to incorporate a wide range of demand distributions, which has value to practitioners, but it does assume identical retailers, which is a significant limitation. Some retailer heterogeneity, however, is easy to incorporate into the model. Specifically, retailer heterogeneity can exist in their cost parameters or in their transportation times from the central warehouse. Since the values of these parameters at one retailer do not influence the lead time distribution at another retailer, incorporating this facet into the model merely requires evaluating the lead time and lead time demand distributions for each distinct retailer type, i.e., XN^{τ} need only be evaluated once.

Heterogeneous retailer batch sizes or demand distributions do create a computational challenge. The results from §2.1 continue to hold, but the evaluation of the retailer's ordering processes, XB^{τ}_{o} and XF^{τ}_{o}, becomes more complex. In particular, with those types of retailer heterogeneity the evaluation of the retailer's ordering processes

Table 4. Policies that minimize total supply chain holding cost while maintaining 99% retailer fill rate.

Scenario			Expected values*									
			Total	Inventory		Backorders		Safety Stock		Fill Rate		Ware. Stockout
Number	R_w	R_r	Cost	Rets.	Ware.	Rets.	Ware.	Rets.	Ware.	Rets. (%)	Ware. (%)	Prob. (%)
1	−1	2	10.40	10.40	0.00	0.00	0.80	6.40	−1.61	99.4	0.0	100
2	−2	2	10.90	10.48	0.43	0.00	0.73	6.47	−2.60	99.4	35.7	100
3	−1	2	16.40	16.40	0.00	0.00	0.80	6.60	−4.55	99.8	0.0	100
4	−1	1	18.40	12.99	5.42	0.00	0.21	3.19	−4.55	99.4	72.2	100
9	−1	2	83.23	83.23	0.00	0.03	6.40	51.21	−6.54	99.4	0.0	100
10	−4	2	81.74	81.74	0.00	0.04	7.90	49.71	−8.92	99.2	0.0	100
11	0	1	103.24	102.46	0.77	0.04	3.17	24.03	−4.96	99.2	31.4	71
12	−1	1	105.24	102.83	2.41	0.04	2.81	24.40	−8.75	99.3	46.9	100
17	9	5	18.04	15.62	2.43	0.04	0.43	11.57	1.93	99.0	89.5	19
18	8	5	18.54	15.66	2.88	0.04	0.38	11.62	1.75	99.1	90.7	20
19	1	5	22.04	20.43	1.61	0.04	1.61	12.32	−1.87	99.0	65.3	45
20	1	5	28.02	21.52	6.50	0.02	0.50	13.43	−1.52	99.6	89.0	41
25	63	5	128.33	125.14	3.19	0.33	3.19	92.82	0.00	99.0	90.0	47
26	61	5	127.84	124.87	2.97	0.34	3.47	92.54	−0.50	99.0	89.2	49
27	17	4	152.32	142.47	9.85	0.32	1.85	77.59	8.00	99.1	94.2	23
28	16	4	154.31	142.68	11.63	0.31	1.63	77.82	9.77	99.1	94.9	21
33	24	6	21.05	18.55	2.50	0.05	1.50	14.50	0.93	99.0	70.7	37
34	23	6	21.54	18.69	2.86	0.04	1.35	14.65	0.75	99.1	73.5	39
35	6	5	26.03	21.22	4.81	0.03	0.81	13.12	2.13	99.4	83.9	24
36	5	5	28.03	21.21	6.83	0.03	0.83	13.11	−1.52	99.3	84.9	45
41	196	5	133.34	124.94	8.41	0.34	3.40	92.60	5.00	99.0	89.4	34
42	195	5	133.84	125.09	8.75	0.33	3.25	92.76	5.50	99.0	89.9	33
43	50	4	156.34	142.09	14.25	0.34	2.25	77.20	12.00	99.0	93.0	22
44	49	4	158.33	142.34	15.99	0.33	1.99	77.46	13.77	99.0	93.8	20
49	7	3	8.02	7.40	0.61	0.04	0.62	6.08	−0.01	99.1	84.5	38
50	6	3	8.51	7.47	1.04	0.03	0.55	6.16	0.17	99.2	86.3	34
51	−1	5	13.99	13.99	0.00	0.01	8.01	7.34	−9.86	99.6	0.0	100
52	−1	4	16.02	14.52	1.50	0.04	3.51	7.84	−9.51	99.0	46.8	100
57	59	3	60.12	59.67	0.45	0.29	4.54	49.11	−4.09	99.1	85.8	79
58	57	3	59.65	59.25	0.40	0.32	4.98	48.66	−4.59	99.0	84.4	81
59	18	2	92.13	79.39	12.74	0.30	0.82	25.98	11.91	99.1	97.4	16
60	17	2	94.12	79.47	14.65	0.30	0.74	26.06	13.69	99.1	97.7	14
65	9	9	34.05	31.19	2.85	0.04	0.85	23.16	1.73	99.1	80.4	25
66	7	9	33.55	31.00	2.55	0.05	1.05	22.96	0.25	99.0	76.4	37
67	2	8	38.04	33.42	4.62	0.04	0.62	22.33	1.91	99.3	86.0	19
68	0	9	40.04	36.27	3.77	0.04	1.76	25.18	−5.77	99.3	66.9	77
73	55	9	248.41	246.93	1.48	0.39	9.47	182.63	−7.99	99.0	70.4	74
74	53	9	247.92	246.54	1.38	0.40	9.87	182.24	−8.49	99.0	69.2	75
75	21	7	264.39	240.07	24.32	0.37	0.31	151.26	24.00	99.0	99.0	5
76	11	8	262.42	260.34	2.09	0.41	12.08	171.49	−10.31	99.0	62.8	69

*Retailer values are totals for all retailers, warehouse values are in units, "Ware. Stockout Prob" is the probability a backorder occurs between the period a system batch is ordered and the period it arrives.

requires knowing *which* retailer ordered before retailer i in the period retailer i places an order, and not merely *how many* retailers ordered before retailer i, as is the case in this model. But this complication is partly an artifact of the allocation method used, namely, the warehouse randomly sorts the retailers each period and fills their order in this sequence, which creates many possible retailer sequences. If the warehouse restricted itself to a reasonable sample of sequences, the computational effort is reduced dramatically.

This research also investigated several heuristics for choosing reorder points. The heuristic which sets warehouse safety stock to zero performs quite well in most cases, but even that heuristic performs poorly in a few cases. Performance is usually poor (but not always) when the warehouse is either restricted to carry no inventory or when it is forced to offer a very high service level to the retailers. The latter result is important because practitioners are often compelled to set very high fill rates at every location in the supply chain. Another approach is to implement the reorder points that are optimal in the continuous inventory review model that provides an approximation for the actual model (which operates with periodic inventory review). Those policies often yield expected costs that are significantly higher than optimal. Overall, it does not appear that a single heuristic will provide good results in all cases, so formal analysis of each case is recommended.

Table 5. Supply chain cost increase when the warehouse reorder point is chosen with a heuristic.

	Percentage Cost Increase over Optimal Policy					Percentage Cost Increase over Optimal Policy			
#	Ware. Carries No Inventory ($R_w = -Q_w$) (%)	Ware. Safety stock= $-Q_w$ (%)	Ware. Safety stock= 0 (%)	At Least 99% Ware. fill rate (%)	#	Ware. Carries No Inventory ($R_w = -Q_w$) (%)	Ware. Safety stock= $-Q_w$ (%)	Ware. Safety stock= 0 (%)	At Least 99% Ware. fill rate (%)
1	14.5	0.0	0.4	28.1	41	65.0	0.4	0.2	8.8
2	32.7	14.1	10.0	23.5	42	65.6	1.5	0.3	9.2
3	0.0	0.0	28.0	66.9	43	50.4	0.8	0.1	9.9
4	27.6	27.6	24.3	50.6	44	52.7	4.9	0.0	10.8
5	0.0	3.0	20.6	67.6	45	66.2	0.7	0.9	14.0
6	33.5	11.6	39.4	60.8	46	66.5	0.1	0.8	14.5
7	0.0	0.0	35.1	88.5	47	48.9	0.5	1.6	15.2
8	19.8	19.8	35.2	69.3	48	50.9	0.3	1.8	16.7
9	36.6	4.0	1.1	5.1	49	6.7	0.0	2.4	9.7
10	37.4	20.3	0.7	5.6	50	25.4	8.3	3.9	12.5
11	1.8	0.0	1.4	14.3	51	0.0	14.2	17.9	49.7
12	15.9	6.6	5.4	14.6	52	14.5	15.2	9.9	52.4
13	8.0	0.5	1.7	15.0	53	10.1	3.6	0.0	32.1
14	12.8	0.7	2.4	16.0	54	26.8	5.6	0.4	33.1
15	4.3	0.0	0.5	17.8	55	0.0	11.1	17.0	71.5
16	16.5	7.8	6.0	18.6	56	12.7	9.5	22.9	82.8
17	14.1	2.6	0.0	20.9	57	27.3	14.0	9.3	0.0
18	16.4	6.0	0.0	22.5	58	30.0	14.4	7.5	0.0
19	2.5	2.7	0.0	19.5	59	5.7	0.8	0.0	11.7
20	15.8	11.6	4.7	35.0	60	16.2	7.8	0.0	12.9
21	12.5	0.0	1.3	35.8	61	25.7	0.2	0.0	5.0
22	17.3	5.3	2.1	37.7	62	29.5	2.2	0.0	5.8
23	0.9	0.7	0.0	35.2	63	4.8	0.1	0.0	17.5
24	13.5	9.3	15.1	57.8	64	11.5	5.8	0.6	19.4
25	27.2	0.2	0.1	4.5	65	11.3	0.5	0.0	21.6
26	27.6	1.6	0.2	4.1	66	12.0	2.7	0.7	22.6
27	18.9	0.7	0.0	7.2	67	6.7	0.0	2.8	24.6
28	21.2	5.7	0.0	8.1	68	15.4	10.6	3.4	35.8
29	23.7	0.9	0.5	5.5	69	11.0	0.2	0.1	38.4
30	24.5	0.5	0.7	4.9	70	11.9	0.4	1.4	40.3
31	14.7	0.5	0.0	9.7	71	5.6	0.0	7.7	43.7
32	17.5	0.4	0.3	11.3	72	15.9	9.6	10.4	62.4
33	32.5	1.0	0.8	29.5	73	24.0	0.8	0.5	3.8
34	36.1	5.1	1.1	31.0	74	24.6	2.2	0.7	4.0
35	22.0	3.6	0.0	27.5	75	19.9	1.3	0.3	5.1
36	28.6	10.7	1.7	40.9	76	22.9	3.6	0.1	5.6
37	33.8	0.6	0.0	52.8	77	27.4	0.0	0.0	8.5
38	35.1	2.2	0.0	54.7	78	27.5	0.6	0.0	8.9
39	22.0	1.4	0.0	48.9	79	21.0	0.0	0.0	10.8
40	26.3	6.9	10.8	68.0	80	24.2	2.6	0.3	11.9

ACKNOWLEDGMENTS

The author would like to thank Morris Cohen, Warren Hausman, Paul Kleindorfer, Peter Liljenberg, Steve Nahmias, David Pyke, Roy Shapiro, Yu-Sheng Zheng, Paul Zipkin and, in particular, Marshall Fisher for their comments on earlier drafts of this paper. The Associate Editor and two referees also provided many useful suggestions to improve this work.

APPENDIX A

SUMMARY OF NOTATION

Constants

Q_r, Q_w Minimum order quantities.

L_r, L_w Fixed transportation times.

N Number of retailers.

μ_r Expected consumer demand at one retailer in one period.

h_r, h_w Per period, per unit inventory holding costs.

p Per period, per unit backorder penalty cost at the retailer.

Decision variables

R_r, R_w Reorder points

Random variables

I_r, I_w On hand inventory.

B_r, B_w Backorder level

IT_r, IT_w Inventory on order.

IP_r, IP_w Inventory position; $IP_j = I_j - B_j + IT_j$, $j \in \{r, w\}$.

Table 6. Supply chain inventory increase when the warehouse reorder point is chosen with a heuristic and retailer fill rate must be at least 99%.

	Percentage Cost Increase over Optimal Policy			
Scenario Number	Ware. Carries No Inventory ($R_w = -Q_w$) (%)	Ware. Safety stock= $-Q_w$ (%)	Ware. Safety stock= 0 (%)	At Least 99% Ware. fill rate (%)
1	0.0	9.6	19.2	38.4
2	18.3	27.5	27.5	36.6
3	0.0	0.0	0.0	24.4
4	43.5	43.5	21.7	43.5
9	0.0	7.2	8.4	14.4
10	0.0	6.1	11.0	17.1
11	27.1	0.0	3.9	15.5
12	19.0	22.8	7.6	15.2
17	33.1	5.5	11.0	22.1
18	21.5	10.8	10.7	21.5
19	18.0	18.0	0.0	18.1
20	14.3	0.1	0.0	14.3
25	49.7	24.0	0.0	9.3
26	49.1	21.7	0.0	9.3
27	36.6	13.0	15.6	7.8
28	31.1	5.2	15.4	7.7
33	52.2	9.4	14.1	28.4
34	41.8	13.9	13.8	27.8
35	15.5	0.1	0.0	15.4
36	28.6	14.3	0.0	28.5
41	92.1	19.3	20.1	13.4
42	90.3	16.3	19.3	13.4
43	53.8	10.2	12.7	10.2
44	68.2	22.6	12.5	10.1
49	49.5	37.0	0.0	37.0
50	23.5	46.7	0.0	34.9
51	0.0	0.2	28.5	28.7
52	24.9	24.9	24.8	49.9
57	59.4	4.8	6.4	14.5
58	58.2	1.6	8.0	16.3
59	21.5	17.2	21.5	8.6
60	46.5	4.2	21.0	8.4
65	17.6	2.9	5.8	23.4
66	14.9	11.9	0.0	26.8
67	0.0	0.0	0.0	10.5
68	10.0	10.0	0.0	20.0
73	28.9	2.8	3.2	10.8
74	28.6	1.6	3.2	11.2
75	15.2	1.5	3.0	0.0
76	25.9	10.6	4.5	1.5

O_r, O_w — Overshoots, $O = R - IP$.

F_r, F_w — Fill rate: The percentage of demand filled immediately from stock.

U_{oj} — The jth batch ordered by a retailer in period t is received by the retailer in period $t + L_r + U_{oj}$, given $O_r = o$ in period t.

D^τ — Consumer demand at a single retailer over τ consecutive periods.

D^τ_{oju} — Consumer demand over periods $[t+1, t+\tau]$, given $O_r = o$ (in period t) and $U_{oj} = u$.

Y^τ_m — Number of batches ordered by m retailers over τ consecutive periods.

XN^τ — Number of batches the retailer orders over periods $[t, t+\tau]$, including only the orders processed after retailer i's period t order.

YB^τ_o — The number of batches retailer i orders over periods $[t-\tau, t-1]$, given that $O_r = o$ in period t.

YF^τ_o — The number of batches retailer i orders over periods $[t+1, t+\tau]$, given that $O_r = o$ in period t.

XB^τ_o — The number of batches the retailers orders over periods $[t-\tau, t]$, including only orders processed before retailer i's period t order.

XF^τ_o — The number of batches the retailers order over periods $[t, t+\tau]$,

Table 7. Comparison of optimal periodic review policies with optimal continuous review policies.

#	Problem Parameters					Periodic Review Optimal Policies			Continuous Review Optimal Policies		% Change in Cost When Using Optimal Continuous Review Policies in a Periodic Review Environment (%)
	μ_r	N	p	Q_r	Q_s	R_s	R_r	Cost	R_s	R_r	
1	0.1	4	20	1	1	0	0	6.23	0	0	0
2	0.1	4	20	1	4	0	0	6.87	−1	0	1
3	0.1	4	20	4	1	−1	0	10.08	−1	0	0
4	0.1	4	20	4	4	−1	0	15.16	−1	−1	9
5	0.1	4	5	1	1	−1	0	4.09	1	−1	36
6	0.1	4	5	1	4	−2	0	4.57	0	−1	35
7	0.1	4	5	4	1	−1	−1	7.28	−1	−1	0
8	0.1	4	5	4	4	−2	−1	11.64	−2	−1	0
9	0.1	32	20	1	1	7	0	41.78	6	0	1
10	0.1	32	20	1	4	6	0	42.08	5	0	1
11	0.1	32	20	4	1	0	0	79.26	2	−1	9
12	0.1	32	20	4	4	−1	0	80.79	1	−1	10
13	0.1	32	5	1	1	4	0	30.27	8	−1	19
14	0.1	32	5	1	4	2	0	30.46	7	−1	19
15	0.1	32	5	4	1	0	−1	55.86	0	−1	0
16	0.1	32	5	4	4	−1	−1	57.19	−1	−1	0
17	1	4	20	1	1	7	4	16.50	8	2	82
18	1	4	20	1	4	6	4	16.69	7	2	78
19	1	4	20	4	1	1	3	20.32	1	2	22
20	1	4	20	4	4	−1	4	22.39	0	2	21
21	1	4	5	1	1	6	3	11.28	8	1	50
22	1	4	5	1	4	5	3	11.48	6	1	54
23	1	4	5	4	1	1	2	14.22	1	0	38
24	1	4	5	4	4	−1	2	15.95	−1	1	18
25	1	32	20	1	1	64	4	118.39	65	2	71
26	1	32	20	1	4	63	4	118.46	63	2	72
27	1	32	20	4	1	15	3	140.09	16	1	57
28	1	32	20	4	4	14	3	140.77	15	1	56
29	1	32	5	1	1	66	2	82.07	63	1	43
30	1	32	5	1	4	64	2	82.14	61	1	44
31	1	32	5	4	1	13	2	100.01	15	0	27
32	1	32	5	4	4	11	2	100.36	13	0	30

including only the orders processed after retailer i's period t order.

Functions and operators

$[a, b]$ The interval from a to b, including a and b.

$\lfloor a/b \rfloor$ The largest integer less than a/b.

$\Pr(A|B)$ The probability of observing event A, given event B.

$E[A]$ Expected value of a random variable A; $E[A] = E_b[E[A|b]]$.

$(a)^+$ This returns the larger of a and zero.

APPENDIX B
THEOREMS

PROOF OF THEOREM 1. Let $IP_i^-(t)$ be location i's inventory position at the beginning of period t, where the warehouse is location 0 and the N retailers are locations 1 through N. Let $IP^-(t) = \{IP_0^-(t), IP_1^-(t), \ldots, IP_N^-(t)\}$. Since demands are independent across time and $\Pr(D^1 = 1) > 0$, $\{IP^-(t)\}$ is an aperiodic irreducible Markov chain. (Given that there is positive probability for unit demand at a retailer, there is positive probability the warehouse demand equals a single batch. Therefore, all of the warehouse's states communicate. Note that $\Pr(D^1 = 1) > 0$ is not a necessary condition; it is merely a relatively innocuous condition which is easy to describe. If $\Pr(D^1 = 1) = 0$, then the demand density must be sufficiently smooth such that all retailer and warehouse inventory positions communicate.) Let $z_0 \in [R_w + 1, R_w + Q_w]$, $z_i \in [R_r + 1, R_r + Q_r]$ for $i > 0$, and $z = \{z_0, z_1, \ldots, z_N\}$.

If

$$\Pr(IP^-(t) = z) = \frac{1}{Q_w}\left(\frac{1}{Q_r}\right)^N \tag{B-1}$$

is the steady state distribution for the Markov chain $\{IP^-(t)\}$, then it follows immediately that $IP_i^-(t)$ is independent of $IP_j^-(t)$, $i \neq j$, the retailer's inventory positions are uniformly distributed on the interval $[R_r + 1, R_r + Q_r]$ and the warehouse's inventory position is uniformly distributed on the interval $[R_w + 1, R_w + Q_w]$.

To demonstrate that (B-1) is the steady state distribution, assume that (B-1) is true for period t. Under that assumption, derive the distribution for $IP^-(t+1)$. If $\Pr(IP^-(t)=z)=\Pr(IP^-(t+1)=z)$, then (B-1) is indeed the steady state distribution. (See Ross 1983, pages 108–109, for details.)

Let

$$\gamma(x,d,r,q)=\left\lfloor \frac{[r-(x-d)+1]^+-1}{q} \right\rfloor+1$$

and

$$\hat{\gamma}(x,d,r,q)=x-d+\gamma(x,d,r,q)q.$$

Note that $\gamma(IP_i^-(t),D^1,R_r,Q_r)$ is the number of batches retailer i orders in period t and $IP_i^-(t+1)=\hat{\gamma}(IP_i^-(t),D^1,R_r,Q_r)$. It holds that

$$\begin{aligned}&\Pr(IP^-(t+1)=z)\\&=\Pr\Bigg(z_0=\hat{\gamma}\Bigg(IP_0^-(t),\sum_{i=1}^{N}\gamma(IP_i^-(t),D^1,R_r,Q_r),R_w,Q_w\Bigg),\\&\qquad z_1=\gamma(IP_1^-(t),D^1,R_r,Q_r),\ldots,\\&\qquad z_N=\gamma(IP_N^-(t),D^1,R_r,Q_r)\Bigg).\end{aligned}\tag{B-2}$$

For a given vector of period t demand, there is a unique $IP_0^-(t)$ such that

$$z_0=\hat{\gamma}\Bigg(IP_0^-(t),\sum_{i=1}^{N}\gamma(IP_i^-(t),D^1,R_r,Q_r),R_w,Q_w\Bigg)$$

and similarly, for each $i>0$ there is a unique $IP_i^-(t)$ such that

$$z_i=\gamma(IP_i^-(t),D^1,R_r,Q_r).$$

According to (B-1), $\Pr(IP_0^-(t)=z_0)=1/Q_w$ and $\Pr(IP_i^-(t)=z_i)=1/Q_r$, $i>0$. Thus, (B-2) simplifies to

$$\Pr(IP^-(t+1)=z)=\frac{1}{Q_w}\left(\frac{1}{Q_r}\right)^N$$

which confirms that (B-1) is indeed the stationary distribution. □

PROOF OF THEOREM 2. The first step is to establish the one-for-one link between the warehouse's inventory position and the sequence in which system batches are shipped. Let $IP_0^-(t)$ be the warehouse's inventory position at the beginning of period t. Number batches in the order they are shipped, with batch 1 being the first batch to be shipped after the beginning of period t. Suppose batch 1 is the vth batch in a system batch. That means that batches $Q_w+1, 2Q_w+1,\ldots$ are also filled with the vth batch within a system batch. (Of course, they are filled from different system batches.) Since $IP_0^-(t)$ is the warehouse's inventory position just before batch 1 is shipped, $IP_0^-(t)$ is also the warehouse's inventory position just before batches $Q_w+1, 2Q_w+1,\ldots$ are shipped. Therefore, whenever the warehouse's inventory position is $IP_0^-(t)$, the next batch shipped is drawn from the vth batch in some system batch. Further, the vth batch in any system batch only fills the subsequent demand when the warehouse's inventory position is $IP_0^-(t)$. To exploit this result, define $z_0(x,j,v)$ to be the warehouse's period t beginning inventory position such that retailer i's jth batch ordered in period t is filled with the vth batch in some system batch given the other retailers order x batches in period t that the warehouse ships before retailer i's period t order. By the previous reasoning, note that $z_0(x,j,v)$ is a function.

From Theorem 1, the locations' inventory positions at the start of period t are independent. Hence, O_r and XB^0 are independent. It follows that

$$\begin{aligned}&\Pr(O_r=o,V=v)=\Pr(O_r=o)\\&\quad\cdot\Bigg(\sum_{x=0}^{\infty}\Pr(XB^0=x)\Pr(IP_0^-(t)=z_0(x,j,v))\Bigg).\end{aligned}\tag{B-3}$$

Since $IP_0^-(t)$ is uniformly distributed on the interval $[R_w+1, R_w+Q_w]$,

$$\Pr(IP_0^-(t)=z_0(x,j,v))=\frac{1}{Q_w}.\tag{B-4}$$

Combining (B-3) and (B-4) yields

$$\Pr(O_r=o,V=v)=\Pr(O_r=o)\frac{1}{Q_w},$$

which demonstrates that Q_r and V are independent and V is uniformly distributed on the interval $[R_w+1, R_w+Q_w]$. □

PROOF OF THEOREM 3. Consider a single retailer that ends τ periods with an inventory position K. From (6), a single retailer orders

$$Y(2R_r+1+Q_r-K,D^\tau)$$

batches over τ periods. Since in steady state K is uniformly distributed on the interval $[R_r+1, R_r+Q_r]$, $2R_r+1+Q_r-K$ is also uniformly distributed on the interval $[R_r+1, R_r+Q_r]$, which means that the retailer orders

$$Y_1^\tau=Y(IP_r^-,D^\tau)=Y(2R_r+1+Q_r-K,D^\tau)$$

batches. So the result holds for one retailer. By convolution, the result holds for Y_m^τ. □

APPENDIX C

EVALUATION OF O_r, O_w, AND IP_o^-

Begin by defining a new random variable, $\Omega=R_r-(IP_r^- - D^1)$, i.e., $R_r-\Omega$ is the retailer's inventory position after demand within a period but before any replenishment has been requested. Let ω be the realization of Ω, $\omega\in[-Q_r,\bar{d}-1]$. A replenishment is requested in a period only when $\Omega\geq 0$, in which case note that the ω is the overshoot. Recall that IP_r^- is the retailer's inventory at the beginning of the period, so in steady state IP_r^- is uniformly distributed on the interval $[R_r+1, R_r+Q_r]$.

If j is the starting inventory position of the retailer and $\Omega=\omega$, then demand in the period is $j-R_r+\omega$. So

$$\Pr(\Omega=\omega) = \frac{1}{Q_r}\sum_{j=R_r+1}^{R_r+Q_r}\Pr(D^1=j-R_r+\omega).$$

$$= \frac{1}{Q_r}(\Pr(D^1\le Q_r+\omega)-\Pr(D^1\le\omega))$$

An order is placed whenever $\Omega\ge 0$, so according to Bayes' Theorem:

$$\Pr(O_r=\omega)=\Pr(\Omega=\omega|\omega\ge 0)$$

$$=\frac{\frac{1}{Q_r}(\Pr(D^1\le Q_r+\omega)-\Pr(D^1\le\omega))}{\frac{1}{Q_r}\sum_{j=R_r+1}^{R_r+Q_r}(1-\Pr(D^1\le j-R_r-1))},$$

which yields

$$\Pr(O_r=o)=\frac{\Pr(D^1\le Q_r+o)-\Pr(D^1\le o)}{\sum_{j=0}^{Q_r-1}(1-\Pr(D^1\le j))}.$$

The maximum overshoot, $\bar{o}_r=\bar{d}-1$.

The evaluation of O_w is analogous to O_r: replace R_r and Q_r with R_w and Q_w; replace D^1 with Y_N^1.

Now consider the evaluation of IP_o^-. Suppose within a period ω is observed. What was the inventory position at the start of the period conditional on this observation? According to Bayes' Theorem,

$$\Pr(IP_r^-=k|\Omega=\omega)=\frac{\Pr(IP_r^-=k \text{ and } \Omega=\omega)}{\Pr(\Omega=\omega)}$$

$$=\frac{\frac{1}{Q_r}\Pr(D^1=k-R_r+\omega)}{\frac{1}{Q_r}(\Pr(D^1\le Q_r+\omega)-\Pr(D^1\le\omega))}.$$

But since $\omega=o$ when $\omega\ge 0$, by definition, $\Pr(IP_o^-=k)=\Pr(IP_r^-=k|\Omega=o)$, i.e.,

$$\Pr(IP_o^-=k)=\frac{\Pr(D^1=k-R_r+o)}{\Pr(D^1\le Q_r+o)-\Pr(D^1\le o)}.$$

APPENDIX D

EVALUATION OF $E[S_{ojuc}]$

For notational clarity, define $\eta^\tau(d)$ as the following series:

$$\eta^\tau(d)=\Pr(D_{oju}^\tau\le d)+\Pr(D_{oju}^{\tau+1}\le d)+\cdots.$$

Hence

$$\eta^{u+L_r+1}(R_r-o+(j-1)Q_r+c-1)=E[S_{ojuc}].$$

Note that for $\tau\ge u+L_r+1$,

$$\eta^\tau(d)=\Pr(D_{oju}^\tau\le d)+P(D^1=0)\sum_{k=\tau}^{\infty}\Pr(D_{oju}^k\le d)$$

$$+\sum_{l=1}^{\min\{d,\bar{d}\}}\Pr(D^1=l)\sum_{k=\tau}^{\infty}\Pr(D_{oju}^k=d-l).$$

The above holds because consumer demands in period $t+u+L_r+1$ and afterwards are independent of u (because once a batch arrives at a retailer, consumer demand afterwards has no influence on its arrival date). Simplification of the above reveals

$$\eta^\tau(d)=\frac{\Pr(D_{oju}^\tau\le d)+\sum_{l=1}^{\min\{d,\bar{d}\}}\Pr(D^1=l)\eta^\tau(d-l),}{\Pr(D^1>0)}. \tag{D-1}$$

The expression (D-1) is recursive and is solved with finite effort.

REFERENCES

Agrawal, N., S. A. Smith. 1996. Estimating negative binomial demand for retail inventory management with unobservable lost sales. *Naval Res. Logist.* **43** 839–861.

Axsäter, S. 1990. Simple solution procedures for a class of two-echelon inventory problems. *Oper. Res.* **38** (1) 64–69.

——. 1993. Exact and approximate evaluation of batch-ordering polices for two-level inventory systems. *Oper. Res.* **41** (4) 777–785.

——. 2000. Exact analysis of continuous review (R, Q)-policies in two-echelon inventory systems with compound Poisson demand. *Oper. Res.* **48** (5) 686–696.

Cachon, G., M. Fisher. 2000. Supply chain inventory management and the value of shared information. *Management Sci.* **46** (8) 1032–1048.

Chen, F., Y.-S. Zheng. 1997. One-warehouse multi-retailer systems with centralized stock information. *Oper. Res.* **45** (2) 275–287.

Cheung, K., W. Hausman. 2000. An exact performance evaluation for the supplier in a two-echelon inventory system. *Oper. Res.* **48** (4) 646–653.

Cohen, A. S., P. Kleindorfer, H. Lee. 1986. Optimal stocking policies for low usage items in multi-echelon inventory systems. *Naval Res. Logist. Quart.* **33** 17–38.

Devermeyer, B., L. 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*, Vol. 16, 163–193.

Graves, S. C. 1985. A multi-echelon inventory model for a repairable item with one-for-one replenishment. *Management Sci.* **31** (10) 1247–1256.

——. 1996. A multiechelon inventory model with fixed replenishment intervals. *Management Sci.* **42** (1) 1–18.

Hausman, W., N. Erkip. 1994. Multi-echelon vs single-echelon inventory control policies. *Management Sci.* **40** (5) 507–602.

Jackson, P. L. 1988. Stock allocation in a two-echelon distribution system or "what to do until your ship comes in." *Management Sci.* **34** (7) 880–895.

Liljenberg, Peter. 1996. The value of centralized information in a two-echelon inventory system with stochastic demand. Working paper, Lund University.

Lee, H., C. Billington. 1993. Material management in decentralized supply chains. *Oper. Res.* **41** (5) 835–847.

——. K. Moinzadeh. 1987a. Operating characteristics of a two-echelon inventory system for repairable and consumable items under batch ordering and shipment policy. *Naval Res. Logist. Quart.* **34** 365–380.

——, ——. 1987b. Two-parameter approximations for multi-echelon repairable inventory models with batch ordering policy. *IIE Trans.* **19** (2) 140–149.

McGavin, E., L. Schwarz, J. Ward. 1993. Two-interval inventory allocation policies in a one-warehouse N-identical retailer distribution system. *Management Sci.* **39** (9) 1092–1107.

Moinzadeh, K., H. Lee. 1986. Batch size and stocking levels in multi-echelon repairable systems. *Management Sci.* **32** (12) 1567–1581.

Muckstadt, J. A., L. J. Thomas. 1980. Are multi-echelon inventory methods worth implementing in systems with low-demand rates? *Management Sci.* **26** (5) 483–494.

Nahmias, S., S. A. Smith. 1994. Optimizing inventory levels in a two-echelon retailer system with partial lost sales. *Management Sci.* **40** (5) 582–596.

Rosenbaum, B. 1981. Service level relationships in a multi-echelon inventory system. *Management Sci.* **27** 926–945.

Ross, S. 1983. *Stochastic Processes*. John Wiley and Sons, New York.

Schwarz, L. B., B. L. Deuermeyer, R. D. Badinelli. 1985. Fill-rate optimization in a one-warehouse N-identical retailer distribution system. *Management Sci.* **31** (4) 488–498.

Sherbrooke, C. G. 1968. METRIC: A multi-echelon technique for recoverable item control. *Oper. Res.* **16** (1) 122–141.

Svoronos, A., P. Zipkin. 1988. Estimating the performance of multi-level inventory systems. *Oper. Res.* **36** (1) 57–72.

Tempelmeier, H. 1993. Safety stock allocation in a 2-echelon distribution system. *Euro. J. Oper. Res.* **65** (1) 96–117.

Veinott, A. F. 1965. Optimal policy for a multi-product, dynamic, nonstationary inventory problem. *Management Sci.* **12** 206–222.