Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Ending Inventory Valuation in Multiperiod Production Scheduling

Marshall Fisher, Kamalini Ramdas, Yu-Sheng Zheng,

Ending Inventory Valuation in Multiperiod Production Scheduling

Marshall Fisher • Kamalini Ramdas • Yu-Sheng Zheng
Operations and Information Management Department, The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104
Darden Graduate School of Business Administration, University of Virginia, Charlottesville, Virginia 22906-6550
Operations and Information Management Department, The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104
fisher@wharton.upenn.edu • ramdask@darden.virginia.edu • zheng@wharton.upenn.edu

When making lot-sizing decisions, managers often use a model horizon T that is much smaller than any reasonable estimate of the firm's future horizon. This is done because forecast accuracy deteriorates rapidly for longer horizons, while computational burden increases. However, what is optimal over the short horizon may be suboptimal over the long run, resulting in errors known as *end-effects*. A common end-effect in lot-sizing models is to set end-of-horizon inventory to zero. This policy can result in excessive setup costs or stockouts in the long run.

We present a method to mitigate end-effects in lot sizing by including a valuation term $V(I_T)$ for end-of-horizon inventory I_T, in the objective function of the short-horizon model. We develop this concept within the classical EOQ modeling framework, and then apply it to the dynamic lot-sizing problem (DLSP). If demand in each period of the DLSP equals the long-run average demand rate, then our procedure induces an optimal ordering policy over the short horizon that coincides with the long-run optimal ordering policy. We test our procedure empirically against the Wagner-Whitin algorithm and the Silver Meal heuristic, under several demand patterns, within a rolling horizon framework. With few exceptions, our approach significantly outperforms the other approaches tested, for modest to long model horizons. We discuss applicability to more general lot-sizing problems.
(*End Effects; Dynamic Lot Sizing; Ending Inventory Valuation*)

1. Introduction

When using lot-sizing models, managers must select an appropriate model horizon T over which to plan production. In practice, the model horizon selected for use is often much smaller than any reasonable estimate of the future horizon faced by the firm. For example, firms routinely make lot-sizing decisions over a quarterly or yearly time horizon, even though they expect to be in business for a much longer period of time. The key reasons for this short-run bias are the impracticality of obtaining accurate demand forecasts for periods that are far in the future, and the increased computational burden associated with longer horizons. Yet the drawback of this approach is that what is optimal in the short run may not be optimal in the long run. This results in *end-effects*, or errors arising from the use of a short model horizon (Grinold 1983). A common end-effect in lot-sizing models is to leave zero ending inventory at the end of the model horizon. While optimal over the short model horizon, this policy can result in excessive setup costs or stockouts in the longer run. Several researchers have addressed the issue of end-effects in lot sizing. Wagner and Whitin (1958) and Federgruen

0025-1909/01/4705/0679$5.00
1526-5501 electronic ISSN

and Tzur (1994) have formulated various sufficient conditions on the values of demand for a particular model horizon T to yield a solution that is optimal for an infinite horizon. Many authors have developed methods described as planning horizon procedures that eliminate the need to forecast unnecessary future data (Zabel 1964, Eppen et al. 1969, Lundin and Morton 1975, Chand et al. 1992). Lee and Denardo (1986) and Chen et al. (1995) have examined the worst-case bound on the error induced by imposing a finite model horizon.

We pursue a different approach for dealing with end-effects in lot sizing. The key conceptual contribution of our paper is the idea that end-effects in lot-sizing models can be mitigated by including a valuation term for end-of-horizon inventory in the objective function of the short horizon model. By explicitly modeling the value to the firm of end-of-horizon inventory, an appropriate level of ending inventory can be induced: We thus provide a rational method to incorporate a longer-term perspective in a model whose horizon is shorter than the "true" problem horizon. Conceptually, our approach builds on the "salvage value" approach developed by Grinold (1983) to reduce end-effects in multistage planning models, in which an "ad hoc value" is placed on resources carried over from the initial transient phase to the subsequent stationary phase of such models.

We introduce this concept and build intuition in the context of the classical EOQ modeling framework: a single item facing constant demand rate D, a setup cost for order placement K, and a holding cost rate h, over a finite horizon of length T. We prove that if inventory I_T left over at the end of the horizon T is assigned value according to the valuation function

$$V(I_T) = K - \frac{h}{2D}(x^* - I_T)^2, \tag{1.1}$$

where x^* is the optimal order quantity for an infinite horizon problem with the same demand and cost parameters as above, then this results in the long-run optimal ordering policy.

Intuitively, $V(I_T)$ represents the future setup cost avoided as a result of having ending inventory I_T on hand. To estimate the impact on future costs of holding I_T at the end of the model horizon T, we assume a constant demand rate D for an infinite horizon beyond T. Then the optimal future policy (given by the well-known EOQ result) is to order $x^* = \sqrt{2KD/h}$ whenever inventory reaches zero, resulting in a long-run optimal cost per unit of time $C^* = \sqrt{2KDh}$. Our ending inventory valuation function $V(I_T)$ can be shown to measure the reduction in setup and holding cost beginning with I_T units compared with the long-run average cost beginning with zero units. The I_T units are depleted in (I_T/D) time units for a total inventory cost of $hI_T^2/2D$. Subtracting this from $(I_T/D)C^*$ (the average cost of supplying I_T units under the optimal EOQ policy) gives a cost reduction of $(I_T/D)C^* - hI_T^2/2D$, which after some algebraic manipulation reduces to the Valuation Function (1.1).

End-of-horizon inventory can be viewed as a storage vehicle for the setup cost. The Valuation Function (1.1) has several desirable properties that entertain an intuitive interpretation. The value of inventory I_T monotonically increases from 0 to K when I_T increases from 0 to x^*, which means that immediately after a setup the entire setup cost K is transformed into the value stored in inventory, and the value decreases as the inventory is depleted. Because the value function is quadratic, the marginal value of I_T decreases as I_T increases, which indicates that the net value of inventory, as a vehicle to store setup costs, diminishes due to the increasing inventory-holding cost. More specifically, the saving of the setup cost is linear in the time duration for which demand is met by existing inventory, hence linear in the inventory. On the other hand, the marginal inventory-carrying cost increases because the more inventory there is, the longer the "last unit" has to be carried. The derivative of the value function $dV/dI_T = -hI_T/D$ up to a constant, indicating that the value of inventory decreases at a rate that is linear in I_T. In fact, the marginal value of any inventory exceeding x^* is negative because the marginal cost of carrying this additional inventory through the period over which it will be used exceeds the marginal saving of the setup cost.

We refine and implement the idea of using ending-inventory valuation to mitigate end-effects in the context of the dynamic lot-sizing problem (DLSP). This problem focuses on finding a cost-minimizing production schedule for a single product with no capacity constraints, given deterministic but variable demand in each period over a model horizon of length T, a fixed cost K per setup, and an inventory holding-cost rate h. Wagner and Whitin (1958) developed an $O(T^2)$ dynamic programming algorithm to solve the T period DLSP. More recently, Federgruen and Tzur (1991), Wagelmans et al. (1992), and Aggarwal and Park (1993) have developed $O(T)$ algorithms for this problem. Because the DLSP ignores all demand beyond the model horizon T, the optimal ending inventory dictated by these algorithms always equals zero. We modify the DLSP by assigning a value to ending inventory according to the end-of-horizon inventory Valuation Function (1.1), where I_T denotes ending inventory in period T, and D is the long-run average demand rate. If the underlying demand distribution is stationary, D can be approximated in practice based on past demand; we will discuss later how for a nonstationary demand distribution, D can be replaced by a dynamically updated long-run demand forecast. The estimate for D would typically come off a company's standard forecasting system and could vary depending on seasonality, forecasted sales growth, etc. Consider first the special case where the demand in each period of the DLSP is constant and equals D. As before, let $x^* = \sqrt{2KD/h}$ denote the classical EOQ based on the above cost parameters and a continuous demand rate equal to D over an infinite horizon. It can be shown that if $x^* = \sqrt{2KD/h}$ covers an integral number of periods of demand in this simplified DLSP, then the use of our valuation function results in an optimal ordering policy over the model horizon T that coincides with the long-run optimal ordering policy (for an infinite horizon DLSP with constant demand D in each time period).

For the general DLSP with time-varying demand, we show that the use of our valuation function, with the demand rate D appropriately chosen, performs very well when compared with existing lot-sizing methods. Because we approximate future demand in each period by the long-term average demand rate D, the data requirements for our approach seem to be a reasonable compromise between totally ignoring demand beyond the model horizon T and requiring detailed period by period demand forecasts far into the future.

We will compare our solution procedure for the DLSP with two popular existing solution procedures: the Wagner-Whitin algorithm and the well-regarded Silver Meal heuristic as enhanced by Blackburn and Millen (1980). Firms typically make production-planning decisions on a rolling horizon basis; they develop a production schedule over a short horizon, implement only the first lot-sizing decision, and then roll the horizon forward to the point when the next decision needs to be made. This is done to use as much and as accurate demand information as is possible while making each lot-sizing decision. We therefore compare our solution procedure with the other two by testing it empirically, using a variety of demand patterns, within a rolling horizon framework. Blackburn and Millen (1980) have shown that in this framework, the optimal Wagner-Whitin algorithm may compare poorly with their enhanced Silver Meal heuristic. We will show that our solution procedure generally performs significantly better than both of these procedures. We will test and discuss the applicability of our solution procedure over stationary demand distributions (uniform and normal), nonstationary demand distributions (seasonal demand with a normal disturbance term, linearly increasing, and linearly decreasing demand), and a distribution with correlated demand. For our empirical analysis, we assume that the true horizon is 300 periods, and use a rolling horizon scheme with model horizons of $T = 2$ to $T = 20$ to compare the cost of our approach relative to the true optimum with that achieved by optimization without ending-inventory valuation and by the enhanced Silver Meal heuristic. For modest to long planning horizons, our approach both outperformed the other two approaches and came consistently close to the true long-run optimum, except in the cases of extreme seasonality, steeply increasing, or steeply decreasing, demand.

The rest of this paper is organized as follows. In the next section, we prove that in the classical EOQ modeling framework, the use of our ending-inventory valuation function $V(I_T)$ results in the long-run optimal ordering policy for a problem with model horizon T.

In §3, we modify the classical Wagner-Whitin algorithm for the DLSP, for the case where end-of-horizon inventory is valued according to $V(I_T)$. In §4, we discuss the empirical tests performed and the results of our analyses. We conclude by discussing the applicability of the concept of using ending-inventory valuation to mitigate end-effects to more general lot-sizing problems.

2. Ending-Inventory Valuation for the Finite Horizon EOQ Model

In this section we consider the finite horizon EOQ model: a single item facing a constant demand rate D, a setup cost for order placement K, and a holding-cost rate h, over a model horizon of length T. When the model horizon is infinite, the optimal policy orders a fixed amount whenever the inventory reaches zero. If we order x every $t = x/D$ time units, we incur a cost per unit time of $K/t + (hD/2)t$, which is minimized at $\tau^* = \sqrt{2K/hD}$ with minimum average cost $C^* = \sqrt{2DKh}$ per unit of time.

When the model horizon T is finite, it is well known (Carr and Howe 1962) that an optimal policy places equal orders when the inventory reaches zero, and leaves no inventory at the end of the horizon. Due to the end-of-horizon effect, the optimal order quantity is a function of T that, in general, differs from $x^* = \tau^* D$, resulting in an average cost exceeding C^*.

We show here that use of the finite horizon model over time with ending-inventory Valuation Function (1.1) provides the infinite horizon optimal solution.

We first establish several properties of an optimal policy for the finite horizon problem with ending-inventory valuation given by (1.1).

LEMMA 1. (*Zero-inventory ordering*) *An optimal policy places orders only when there is no inventory.*

PROOF. If an order is placed before the previous order is completely depleted, a cheaper solution can always be obtained by increasing the size of the order placed by the amount of inventory on hand and reducing the size of the previous order by the same amount.

With the zero-inventory ordering restriction, a policy can be specified by its replenishment epochs T_i, $i = 1, \ldots, n\,(0 = T_0 < T_1 < \cdots < T_n < T_{n+1} = T)$ along with the size of the last order placed at T_n. For given replenishment epochs $T_0, T_1, \ldots, T_n$, let τ_i be the length of the ith replenishment interval, i.e., $\tau_i = T_i - T_{i-1}$, $i = 1, \ldots, n+1$. It is easy to show that the total cost for each of the first n cycles is at least $C^*\tau_i$. □

LEMMA 2. *For any order interval (T_{i-1}, T_i), $i = 1, \ldots, n$, with zero ending inventory, the cost incurred in this interval is greater than or equal to $C^*\tau_i$. The equality holds only when $\tau_i = \tau^*$.*

PROOF. The cost per unit time for the interval (T_{i-1}, T_i) is given by the strictly convex function $K/\tau_i + (hD/2)\tau_i$, which is minimized at $\tau_i = \tau^*$ with value C^*. □

LEMMA 3. *For any policy with $\tau_{n+1} \le \tau^*$, the optimal order quantity at time T_n is x^*, and the optimal cost incurred in the interval (T_n, T) is $C^*\tau_{n+1}$.*

PROOF. The cost incurred in (T_n, T) as a function of the order quantity x is

$$\begin{aligned} C(x) &= K + \frac{h\tau_{n+1}}{2}(2x - D\tau_{n+1}) - V(x - D\tau_{n+1}) \\ &= \frac{h\tau_{n+1}}{2}(2x - D\tau_{n+1}) + \frac{h}{2D}(x^* - x + D\tau_{n+1})^2. \end{aligned}$$

$C(x)$ is quadratic in x, and its minimum is reached at $x = x^*$ with $C(x^*) = h\tau_{n+1}x^* = C^*\tau_{n+1}$. □

LEMMA 4. *In an optimal order policy, the last interval length $\tau_{n+1} \le \tau^*$.*

PROOF. If $\tau_{n+1} > \tau^*$, the minimum cost for this interval is reached at $x = D\tau_{n+1}$ with final inventory equal to zero (for any positive final inventory, the cost of holding this inventory less its value is always positive if $\tau_{n+1} > \tau^*$). The resulting cost $K/\tau_{n+1} + (hD/2)_{\tau_{n+1}}$ can be strictly reduced by splitting the interval $[T_n, T]$ into $l(\ge 2)$ intervals with the length of the first $l-1$ intervals being τ^* and the length of the last interval being at most τ^*, because by Lemma 3, the order quantity for the last interval is also x^*.

Theorem 1. *For a finite horizon model with the Valuation Function* (1.1) *for end-of-horizon inventory, an optimal policy orders x^* units whenever the inventory reaches zero.*

Proof. In view of Lemmas 2, 3, and 4, for any given policy with reorder intervals $\tau_1, \tau_2, \ldots, \tau_{n+1}$, the cost incurred is

$$C(T) \geq \sum_{i=1}^{n+1} C^* \tau_i = C^* T.$$

The inequality holds with equality only when $\tau_i = \tau^*$, $i = 1, \ldots, n$, $\tau_{n+1} \leq \tau^*$ and x^* units are ordered at time T_i. □

It can be proved, similarly, that for the DLSP with constant demand D in each period, setup cost K, and holding-cost rate h, the use of the Valuation Function (1.1) results in an optimal policy over a T period model horizon that coincides with the long-term optimal policy, provided that $x^* = \sqrt{2KD/h}$ covers an integral number of periods of demand. Even in the absence of this last restriction, a slightly modified value function will provide the same result. This provides the basis of our use of ending-inventory valuation with the DLSP.

3. The Dynamic Lot-Sizing Model with Time Varying Demand

In general, demand in each period is typically not constant. We modify the DLSP by assigning a value to ending inventory according to (1.1) with D representing the long-term average demand rate. Let x_t and I_t denote the amount produced and the ending inventory, respectively, in period t, $t = 1, \ldots, T$. The modified dynamic lot-sizing problem is to determine a production schedule that minimizes setup and holding costs over the model horizon T, less the value $V(I_T)$ of ending inventory I_T. Below we show that the classical Wagner-Whitin algorithm can be modified to solve this problem.

For simplicity, assume that the beginning inventory $I_o = 0$. If $I_{s-1} > 0$ and $x_s > 0$, for some period s in a given schedule, then a schedule that is no more costly can be obtained by reducing the last order before period s by I_{s-1} and setting $x_s = x_s + I_{s-1}$. Hence, in any period of an optimal schedule, production takes place only if inventory at the beginning of the period is zero, i.e., $I_{t-1} x_t = 0$, $t = 1, \ldots, T$. This property implies that for each period t, the optimal x_t must be zero, a cumulative sum of demands in periods t through s for some $s \leq T$ or the cumulative sum of demands in periods t through T plus some ending inventory.

For $1 \leq t < s$, let $C_{t,s}$ denote the minimum cost incurred in periods $t, t+1, \ldots, s-1$, given that production takes place in period t for periods t, $t+1, \ldots, s-1$ for $s \leq T$.

$$C_{t,s} = K + \sum_{j=t+1}^{s-1} (j-t) h d_j. \tag{3.1}$$

For $s = T+1$, the optimal production quantity x_t^{opt} should minimize $C_{t,T+1}(x_t)$, the cost incurred in periods $t, \ldots, T$, less the value of any ending inventory I_T given that production takes place only in period t.

$$\begin{aligned} C_{t,T+1}(x_t) &= K + \sum_{j=t}^{T} h I_j - \left\{ K - \frac{h}{2D}(x^* - I_T)^2 \right\} \\ &= \sum_{j=t}^{T} h\left(x_t - \sum_{i=t}^{j} d_i\right) \\ &\quad + \frac{h}{2D}\left(x^* - x_t + \sum_{i=t}^{T} d_i\right)^2. \end{aligned} \tag{3.2}$$

For each t, $C_{t,T+1}(x_t)$ is quadratic in x_t and is minimized by setting

$$x_t^{\text{opt}} = \begin{cases} \sum_{i=t}^{T} d_i + x^* \\ \quad -(T+1-t)D & x^* > (T+1-t)D \\ \sum_{i=t}^{T} d_i & x^* \leq (T+1-t)D. \end{cases} \tag{3.3}$$

Let f_t denote the total cost less value of ending inventory of a schedule that is optimal for periods t through T, $1 \leq t \leq T$. Define $f_{T+1} = 0$. For each period $t \leq T$, f_t is obtained by the recursion

$$f_t = \min_{t < s \leq T+1} (C_{t,s} + f_s) = C_{t,s(t)} + f_{s(t)}, \tag{3.4}$$

where $C_{t,T+1} = C_{t,T+1}(x_t^{\text{opt}})$. The quantity f_1 is the total cost less value of ending inventory for a schedule that is optimal over the horizon T. The optimal sequence of production periods is 1, $s(1)$, $s(s(1))$, etc. In this sequence, if production occurs in period i, the optimal production quantity in period i is

$$x_i = \begin{cases} \sum_{k=i}^{s(i)-1} d_k & s(i) < T+1 \\ \sum_{k=i}^{T} d_k & s(i) = T+1, \\ & \quad x^* \le (T+1-i)D \\ \sum_{k=i}^{T} d_k + x^* & \\ \quad -(T+1-i)D & s(i) = T+1, \\ & \quad x^* > (T+1-i)D. \end{cases} \tag{3.5}$$

4. Comparison of Methods for Solving the DLSP within a Rolling Horizon Framework

In this section we compare our Ending-Inventory Valuation algorithm (EIV) with the classical Wagner-Whitin algorithm (WW) and the Silver Meal heuristic (SM) as enhanced by Blackburn and Millen.

Silver's and Meal's (1973) heuristic for the T period DLSP works as follows. Let $A[t, s]$ denote the average cost in periods t through s if we set up and produce in period t for periods t through s. Because we assume zero starting inventory, note that a setup occurs in Period 1. Let t_1 denote the smallest $t \le T$ for which $A[1, t+1] > A[1, t]$. Then the order quantity in Period 1 covers demands in periods $1, \ldots, t_1$. Let t_2 denote the smallest period beyond t_1 ($t_1 < t_2 \le T$) with nonzero demand. The next setup occurs in t_2. The process is repeated until the model horizon is reached.

Blackburn and Millen (1980) have pointed out that if demand can be zero in some periods of the DLSP, it is more logical to define $A[t, s]$ to be the average cost in periods t through $s+h$ where h is the number of consecutive zero-demand periods immediately following s. This modification has no impact unless demand is zero, in which case it generally improves performance, so in our computations, we use the variant of the Silver Meal heuristic proposed by Blackburn and Millen.

Because the decision of whether to include the demand for any period in the preceding batch is made without examining demands beyond that period, the performance of the Silver Meal heuristic is insensitive to the length of the model horizon T as long as T exceeds the number of periods covered by the first production lot. Let t^* denote the optimal reorder interval for an infinite horizon DLSP with demand in each period equal to the long-run average demand rate. Because the average reorder interval in the DLSP is t^*, the performance of the myopic Silver Meal heuristic is likely to improve as the length of the model horizon T is increased up to t^*, and is likely to be unaffected by further increases in T.

A rolling horizon scheme solves the T period problem, using either a heuristic or an optimal algorithm, and implements only the first production lot size. The start of the model horizon is then rolled forward to the first future period not covered by this production lot, and the procedure is repeated. In using a rolling horizon scheme to compare our EIV algorithm with the WW algorithm and the SM heuristic, our goal was to measure the distortion relative to the long-run optimum caused by the use of a short model horizon for each lot-sizing method. To obtain a reasonable surrogate for the long-run optimal value, we generated the "actual" demand for 300 periods and computed the optimal 300-period solution. We computed the percentage increase over the optimal cost, when using each lot-sizing method in conjunction with the rolling scheme over the time span of 300 periods.

The experimental setting we used was similar to that of Blackburn and Millen (1980). Demand d_t in each period t, for $1 \le t \le 300$, was generated according to a variety of distributions. For each distribution, we examined the effect of the length of the model horizon, setup and holding costs, and variance of the demand distribution on the performance of the lot-sizing methods.

We assume that in practice it is not possible to have a perfect demand forecast for a 300-period horizon, so "actual" demand over this period would not be known in advance. Therefore, simply using the Wagner-Whitin algorithm over this long horizon is not a viable option in practice. In using our EIV algorithm, we assume that we have a perfect forecast up to the model horizon, beyond which we have a very rough forecast that is adequate to estimate the long-run average demand rate, D. We outline a procedure to estimate D for each type of demand distribution included in our analysis.

Stationary Demand

In this category, we considered two distributions: uniform and normal. The mean μ for both distributions was set at 100, and the range R of the uniform distribution or standard deviation σ of the normal distribution was set to obtain a range of equivalent variances for both distributions. The values used were $R = 0, 35, 75, 150$ and $\sigma = 0, 10, 22, 43$. We set the long-term average demand rate $D = \mu$. In practice, D would be estimated using an approach such as a moving average of historical demand. Depending on the level of estimation error, D might not exactly equal μ. Following Blackburn and Millen (1980), for each demand scenario, the holding cost h was normalized at one and the setup cost K was chosen to yield economic order cycles $t = \sqrt{2K/Dh}$ of length three, four, five, or six periods. With mean demand $D = 100$ and $h = 1$, this yields $K = 50t^2$, for $t = 3, \ldots, 6$.

For each choice of input parameters, we generated eight problems, each consisting of 300 periods of simulated demand. For each problem, the optimal cost was computed over the horizon of 300 periods. Next, the percentage above optimal cost was computed for each of the three lot-sizing methods, in a rolling horizon scheme with the length of the short model horizon T ranging from 2 to 20 periods. The results for normally distributed demand are in Table 1 (we obtained similar results for uniformly distributed demand). We present only the scenarios in which $K = 800$, because the results were similar for all values of setup cost. For each demand scenario, the figures represent average performance of the lot-sizing methods over eight test problems.

Note that use of our ending-inventory valuation formula produced superior results to either the WW algorithm or the SM heuristic except for 8 of the 16 cases in which the model horizon T was 3 or 4. In these 8 cases, all three methods usually incurred a setup only in the first period of the planning horizon, but our ending-inventory valuation induced an ending inventory with $0 < I_T < d_{T+1}$. Hence, the cost of carrying I_T was incurred without a compensating reduction in setup cost because I_T was insufficient to satisfy demand in period $T+1$. This does not occur for $T = 2$ because in this case, with $x^* = 400$ and $D = 100$, our valuation function results $I_T = x^* - 2D = 200$ units, which is typically enough to cover demand in at least one period beyond T,

Table 1 **Percentage Deviation from Optimality for Normally Distributed Demand**

Model	$\sigma = 0$			$\sigma = 10$			$\sigma = 22$			$\sigma = 43$		
Horizon	WW	SM	EIV	WW	SM	EIV	WW	SM	EIV	WW	SM	EIV
2	28.57	28.57	0.00	29.29	29.29	15.03	31.43	31.43	17.14	36.14	36.14	21.00
3	4.76	4.76	0.00	5.24	5.24	15.31	6.85	6.85	17.59	11.15	11.22	20.84
4	0.00	0.00	0.00	0.57	0.83	0.80	2.25	1.85	1.72	6.04	4.11	3.76
5	0.00	0.00	0.00	3.32	0.69	0.61	4.65	1.13	0.90	6.82	1.97	1.93
6	0.00	0.00	0.00	5.20	0.70	0.47	5.11	1.09	0.65	5.31	1.63	0.71
7	0.00	0.00	0.00	1.37	0.70	0.46	1.63	1.09	0.63	2.21	1.60	0.68
8	0.00	0.00	0.00	0.51	0.70	0.32	0.94	1.09	0.44	1.46	1.60	0.46
9	0.00	0.00	0.00	1.14	0.70	0.29	1.29	1.09	0.34	0.78	1.60	0.42
10	0.00	0.00	0.00	1.85	0.70	0.31	1.31	1.09	0.34	0.97	1.60	0.28
11	0.00	0.00	0.00	0.78	0.70	0.22	0.64	1.09	0.26	0.57	1.60	0.16
12	0.00	0.00	0.00	0.42	0.70	0.16	0.42	1.09	0.21	0.38	1.60	0.16
13	0.00	0.00	0.00	0.61	0.70	0.13	0.39	1.09	0.12	0.29	1.60	0.10
14	0.00	0.00	0.00	0.83	0.70	0.12	0.39	1.09	0.15	0.18	1.60	0.11
15	0.00	0.00	0.00	0.44	0.70	0.13	0.30	1.09	0.11	0.18	1.60	0.11
16	0.00	0.00	0.00	0.32	0.70	0.14	0.19	1.09	0.07	0.12	1.60	0.04
17	0.00	0.00	0.00	0.34	0.70	0.13	0.17	1.09	0.09	0.16	1.60	0.04
18	0.00	0.00	0.00	0.52	0.70	0.10	0.18	1.09	0.05	0.10	1.60	0.03
19	0.00	0.00	0.00	0.35	0.70	0.07	0.11	1.09	0.05	0.06	1.60	0.02
20	0.00	0.00	0.00	0.16	0.70	0.08	0.11	1.09	0.05	0.04	1.60	0.02

Note. $\mu = 100$, $K = 800$, $h = 1$ (WW: Wagner Whitin, SM: Silver Meal, EIV: Ending Inventory Valuation).

and therefore reduces the total setup costs over the 300-period horizon. The results for $T = 3$ and $T = 4$ are due to round-off error resulting from the periodic review that is inherent in the DLSP. If it is possible in practice to reduce the length of each review period, we can reduce this round-off error by using finer periods. Using finer periods increases the number of periods, but this is not a problem in terms of computational burden when $T = 3$ or 4, because the number of periods is so small.

Table 2 shows the minimum length of model horizon needed to achieve a fixed deviation from the long-term optimum over 300 periods for the WW and EIV algorithms and the SM heuristic. We considered deviations from the optimum of 1% and 0.5%. The results indicate that for both the normal and uniform distributions, the length of model horizon required to achieve a given deviation is much less (roughly by a factor of 1.5 to 2) by using the ending-inventory valuation principle than by ignoring it. Table 2 also contains similar comparisons for the other demand distributions studied.

A seemingly counterintuitive pattern observed in Table 1 is that for long model horizons, the performance of the WW and EIV results generally improves as variance increases in the case of both uniform and normal distributions. In contrast, the performance of the SM heuristic worsens as variance increases. Blackburn and Millen (1980) also reported that the WW algorithm performs better as variance increases, while the SM heuristic performs worse. This improvement in the performance of the WW and EIV algorithms relative to the SM heuristic with longer horizons and more variability occurs because both these algorithms, which optimize production over the entire planning horizon T, are able to fine-tune ordering decisions so as to exploit variability in forecasted demand.

Table 2 **Minimum Length of Model Horizon Required to Achieve a Given Deviation from Optimum**

	WW	SM	EIV	WW	SM	EIV	WW	SM	EIV	WW	SM	EIV
Uniform		$R = 0$			$R = 35$			$R = 75$			$R = 150$	
1% deviation	4	4	2	11	4	4	11	> 20	5	10	> 20	7
0.5% deviation	4	4	2	15	> 20	7	12	> 20	8	11	> 20	8
Normal		sigma = 0			sigma = 10			sigma = 22			sigma = 43	
1% deviation	4	4	2	11	4	4	11	> 20	5	9	> 20	6
0.5% deviation	4	4	2	19	> 20	6	12	> 20	8	12	> 20	8
Seasonal (M = 1)		$a = 20$			$a = 40$			$a = 60$			$a = 80$	
1% deviation	11	4	5	12	> 20	10	11	> 20	10	10	> 20	10
0.5% deviation	15	> 20	10	15	> 20	11	14	> 20	11	16	> 20	11
Seasonal (M = 4)		$a = 20$			$a = 40$			$a = 60$			$a = 80$	
1% deviation	11	4	5	12	> 20	10	11	> 20	10	10	> 20	9
0.5% deviation	15	> 20	9	15	> 20	11	14	> 20	11	16	> 20	11
Linearly Increasing Trend		$c = 1$			$c = 10$			$c = 20$			$c = 40$	
1% deviation	10	3	4	2	2	3	2	2	2	2	2	2
0.5% deviation	20	3	4	2	2	3	2	2	3	2	2	2
Linearly Decreasing Trend		$c = 1$			$c = 10$			$c = 20$			$c = 40$	
1% deviation	8	3	4	5	2	3	2	2	2	2	2	2
0.5% deviation	12	4	4	12	2	3	8	2	3	2	2	2
Correlated (Myopic Forecast)		sigma = 0			sigma = 5			sigma = 10			sigma = 15	
1% deviation	11	5	8	11	5	7	11	> 20	7	11	> 20	7
0.5% deviation	13	> 20	9	14	> 20	10	13	> 20	10	13	> 20	8
Correlated (Long-Term Forecast)		sigma = 0			sigma = 5			sigma = 10			sigma = 15	
1% deviation	11	6	6	11	> 20	6	10	> 20	7	11	> 20	7
0.5% deviation	13	> 20	9	14	> 20	8	13	> 20	8	12	> 20	8

Note. $\mu = 100$, $K = 800$, $h = 1$ (WW: Wagner Whitin, SM: Silver Meal, EIV: Ending Inventory Valuation).

Linearly Increasing or Linearly Decreasing Demand We generated a demand distribution with linearly increasing trend by adding a linear trend term to a normally distributed random variable. Demand was generated using the formula $d_t = \mu + \sigma z_t + c(t-1)$, for $1 \leq t \leq 300$, where z_t is a standard normal variable and c is an additive trend factor. To obtain a comparable demand distribution d'_t with linearly decreasing trend, we set $d'_t = d_{300-t+1}$, for $1 \leq t \leq 300$. In both cases, we set $\mu = 100$, $\sigma = 10$, $K = 800$, $h = 1$, and $c = 1, 10, 20$, and 40. The results for a linearly increasing trend are in Table 3a, and those for a linearly decreasing trend are in Table 3b. Each entry in these tables is obtained by averaging the results over eight demand simulations.

In this case, demand is either growing or declining steadily over time, as is the average demand rate. We therefore dynamically updated the demand forecast, with $D_t = \mu + c(t-1)$ for growing demand, where D_t denotes the average demand in period t. Similarly, for declining demand, $D_t = \mu + c(T-t)$. In practice, D_t may be calculated using a moving average or a more sophisticated forecasting method that incorporates the trend component. To calculate end-of-horizon inventory for a model horizon of length T that starts in period t, we replaced D with D_{t+T} and x^* with x^*_{t+T}, which is adjusted based on D_{t+T}. The number of periods that will be covered by end-of-horizon inventory is a function of t, T, and the trend component c. For an increasing trend, the number of periods covered by end-of-horizon inventory will decrease steadily over time. Finally, when the demand is large enough, there will be a setup in every period, with no inventory carried forward. For a decreasing trend, the reverse will occur.

We see in Table 3a that when demand exhibits a linearly increasing trend, all three methods tested perform well, and the EIV algorithm performs better than the WW algorithm for most lengths of planning horizon, especially for a small positive trend ($c = 1, 10$). For any particular length of planning horizon, the performance of all three methods appears to improve as the positive trend becomes steeper. As the magnitude of the trend factor increases, the first order quantity is likely to cover fewer and fewer periods of

Table 3a Percentage Deviation from Optimality for a Demand Distribution with Linearly Increasing Trend

	$c=1$			$c=10$			$c=20$			$c=40$		
Model Horizon	WW	SM	EIV	WW	SM	EIV	WW	SM	EIV	WW	SM	EIV
2	5.00	5.00	18.51	0.32	0.32	1.21	0.15	0.15	0.62	0.06	0.06	0.27
3	3.28	0.35	2.98	0.38	0.01	0.13	0.16	0.02	0.04	0.02	0.00	0.01
4	1.66	0.16	0.16	0.13	0.01	0.03	0.07	0.02	0.01	0.06	0.00	0.00
5	2.90	0.15	0.16	0.28	0.01	0.01	0.07	0.02	0.01	0.02	0.00	0.02
6	0.40	0.15	0.16	0.02	0.01	0.02	0.02	0.02	0.01	0.00	0.00	0.00
7	1.97	0.15	0.16	0.26	0.01	0.02	0.09	0.02	0.03	0.02	0.00	0.02
8	0.46	0.15	0.14	0.02	0.01	0.03	0.02	0.02	0.02	0.00	0.00	0.00
9	1.05	0.15	0.13	0.15	0.01	0.01	0.05	0.02	0.01	0.02	0.00	0.02
10	0.50	0.15	0.10	0.03	0.01	0.01	0.01	0.02	0.01	0.00	0.00	0.00
11	0.94	0.15	0.11	0.13	0.01	0.01	0.04	0.02	0.01	0.02	0.00	0.01
12	0.16	0.15	0.07	0.01	0.01	0.01	0.01	0.02	0.01	0.00	0.00	0.00
13	0.94	0.15	0.04	0.12	0.01	0.02	0.04	0.02	0.01	0.01	0.00	0.01
14	0.16	0.15	0.05	0.02	0.01	0.02	0.01	0.02	0.01	0.00	0.00	0.00
15	0.61	0.15	0.08	0.08	0.01	0.01	0.04	0.02	0.01	0.00	0.00	0.00
16	0.21	0.15	0.07	0.01	0.01	0.01	0.01	0.02	0.01	0.00	0.00	0.00
17	0.60	0.15	0.09	0.08	0.01	0.01	0.03	0.02	0.00	0.00	0.00	0.00
18	0.10	0.15	0.05	0.01	0.01	0.01	0.01	0.02	0.01	0.00	0.00	0.00
19	0.59	0.15	0.07	0.06	0.01	0.01	0.01	0.02	0.00	0.00	0.00	0.00
20	0.09	0.15	0.06	0.01	0.01	0.00	0.01	0.02	0.01	0.00	0.00	0.00

Note. $\mu = 100$, $\sigma = 10$, $K = 800$, $h = 1$.

Table 3b Percentage Deviation from Optimality for a Demand Distribution with Linearly Decreasing Trend

	$c = 1$			$c = 10$			$c = 20$			$c = 40$		
Model Horizon	WW	SM	EIV	WW	SM	EIV	WW	SM	EIV	WW	SM	EIV
2	5.12	5.12	18.72	0.45	0.45	1.52	0.28	0.28	0.94	0.28	0.28	0.33
3	3.48	0.54	3.32	1.25	0.09	0.16	0.89	0.17	0.17	0.54	0.20	0.17
4	1.76	0.32	0.26	0.18	0.11	0.19	0.07	0.09	0.12	0.21	0.01	0.23
5	3.14	0.24	0.26	1.04	0.11	0.10	0.65	0.05	0.08	0.25	0.01	0.04
6	0.39	0.24	0.16	0.04	0.11	0.04	0.04	0.05	0.06	0.02	0.01	0.09
7	1.58	0.24	0.11	0.62	0.11	0.03	0.52	0.05	0.04	0.17	0.01	0.02
8	0.49	0.24	0.07	0.11	0.11	0.09	0.03	0.05	0.05	0.01	0.01	0.02
9	0.68	0.24	0.06	0.55	0.11	0.03	0.41	0.05	0.03	0.09	0.01	0.03
10	0.45	0.24	0.06	0.03	0.11	0.02	0.02	0.05	0.04	0.01	0.01	0.00
11	0.64	0.24	0.06	0.51	0.11	0.03	0.32	0.05	0.01	0.04	0.01	0.04
12	0.15	0.24	0.06	0.03	0.11	0.02	0.01	0.05	0.02	0.02	0.01	0.02
13	0.49	0.24	0.06	0.44	0.11	0.01	0.25	0.05	0.01	0.01	0.01	0.01
14	0.16	0.24	0.05	0.03	0.11	0.03	0.01	0.05	0.02	0.02	0.01	0.02
15	0.27	0.24	0.03	0.41	0.11	0.02	0.19	0.05	0.01	0.00	0.01	0.00
16	0.20	0.24	0.02	0.01	0.11	0.01	0.01	0.05	0.02	0.02	0.01	0.02
17	0.26	0.24	0.02	0.37	0.11	0.01	0.14	0.05	0.00	0.01	0.01	0.01
18	0.09	0.24	0.02	0.02	0.11	0.01	0.01	0.05	0.02	0.00	0.01	0.00
19	0.29	0.24	0.02	0.36	0.11	0.02	0.10	0.05	0.00	0.00	0.01	0.00
20	0.05	0.24	0.03	0.01	0.11	0.01	0.01	0.05	0.01	0.00	0.01	0.00

Note. $\mu = 100$, $\sigma = 10$, $K = 800$, $h = 1$.

demand, so that it becomes less likely that the end-of-horizon constraints will affect the first order quantity. When demand exhibits a linearly increasing trend, the SM heuristic is also a good choice. It is easy to compute and outperforms both WW and EIV for two of the four values of the trend parameter c.

When demand exhibits a relatively slow linear decline ($c = 1$ and $c = 10$, in Table 3b), EIV outperforms both WW and SM, for a model horizon of four or more periods. Interestingly, we find that the SM heuristic performs worse for decreasing trend than for increasing trend for all values of c. For the SM heuristic, the quantity produced in a production period includes demand for subsequent periods so long as average cost is decreasing. When including the demand for a period in the previous batch increases average cost, a new batch is started in that period. However, with a decreasing trend, the probability that it is in fact optimal to include the demand for this and subsequent periods in the previous production batch is higher, relative to the case of increasing trend. As in the case of increasing trend, we find that for any particular length of planning horizon, the performance of all three methods improves for more steeply decreasing trend. This is because for more steeply decreasing trend, periods early in the 300-period horizon are more likely to have demand large enough to incur a setup. For more extreme values of trend, EIV does not strictly dominate the other heuristics, but the differences are small.

Seasonal Demand

For this case, we generated demand according to the formula used by Baker et al. (1978):

$$d_t = \mu + \sigma z_t + a \sin\left[\frac{2\pi}{b}(t + b/4)\right],$$

where z_t is a standard normal variable generated by Monte Carlo simulation, and a and b denote the amplitude of the seasonal component and the length of the seasonal cycle in periods, respectively. For a seasonal demand pattern, the long-run average demand rate D (which can be approximated by μ) will not provide a good estimate of the demand pattern during different seasons. To account for this, we introduce a seasonally adjusted demand estimate D_t.

For any period t, define

$$D_t = \mu + a\sin\left[\frac{2\pi}{b}(t+b/4)\right].$$

Note that D_t would be estimable in practice, because companies know the seasonality of their demand and incorporate this into their demand forecasts. Further, let $D_{t,T}$ denote the demand estimate used in calculating end-of-horizon inventory, for a model horizon of length T that starts in period t, $1 \le t \le 300$. We define

$$D_{t,T} = \frac{1}{t^*}\sum_{i=t+T}^{i=t+T+t^*-1} D_t,$$

where, as before, t^* is the optimal number of periods between orders assuming a constant demand D in each period.

To calculate end-of-horizon inventory for a model horizon of length T that starts in period t, we replace D with $D_{t,T}$ and x^* with $x^*_{t,T}$, where the latter is seasonally adjusted based on $D_{t,T}$. With these adjustments, $D_{t,T}$ and $x^*_{t,T}$ will be higher in the case of a seasonal demand peak than in the case of a seasonal trough at the end of the model horizon. In practice, $D_{t,T}$ would be based on a seasonally adjusted demand forecast for the t^* periods including and immediately following $t+T$.

To compare the solution procedures, we set $\mu = 100$ and $\sigma = 10$, $K = 800$, and $h = 1$, and varied a and b. We set $a = 20, 40, 60, 80$, and $b = 4, 12$, and 52 to reflect quarterly, monthly, and weekly demand data. Because the results were similar for all values of b, we report on only the case where $b = 12$ in Table 4. Each entry in this table is obtained by averaging the results over eight demand simulations for each set of input parameters. We found that for a model horizon of five or more periods, the EIV algorithm typically performs significantly better than the SM and the WW methods if the degree of seasonality is not extreme. In the case of extreme seasonality ($a = 60$ or $a = 80$), the WW algorithm occasionally outperforms the EIV algorithm even for longer model horizons, but the differences are small. For long model horizons, an improvement in performance for the WW and the EIV algorithms when variability (in terms of seasonality) increases occurs less consistently than in the stationary case. This may be because the degree of variability considered in this case is much higher

Table 4 Percentage Deviation from Optimality for a Seasonal Demand Distribution

Model	$a = 20$			$a = 40$			$a = 60$			$a = 80$		
Horizon	WW	SM	EIV	WW	SM	EIV	WW	SM	EIV	WW	SM	EIV
2	29.70	29.70	16.93	31.22	31.22	16.15	34.84	34.84	15.91	40.68	40.68	13.63
3	5.57	5.57	16.98	6.81	6.81	17.19	9.76	9.76	16.22	14.54	14.54	12.14
4	0.89	0.99	4.02	2.07	1.87	6.43	4.89	4.37	5.08	9.39	8.13	6.89
5	3.59	0.68	0.82	3.96	1.40	1.13	6.55	3.94	2.18	9.69	7.73	5.03
6	4.40	0.74	0.72	2.42	1.48	1.42	1.85	2.73	1.00	2.75	2.64	3.37
7	1.79	0.74	0.73	2.47	1.55	1.02	2.53	2.93	1.57	2.11	4.23	2.13
8	0.68	0.74	0.58	1.04	1.58	0.97	2.24	3.29	1.32	2.99	3.73	2.06
9	1.18	0.74	0.36	1.13	1.58	1.06	1.81	3.54	1.44	1.72	4.82	0.93
10	1.79	0.74	0.29	1.77	1.58	0.57	1.03	3.54	0.72	0.87	4.80	0.52
11	0.93	0.74	0.24	1.03	1.58	0.25	0.88	3.54	0.23	0.46	4.80	0.19
12	0.30	0.74	0.19	0.29	1.58	0.18	0.23	3.54	0.16	0.22	4.80	0.26
13	0.51	0.74	0.14	0.58	1.58	0.13	0.51	3.54	0.15	0.37	4.80	0.11
14	0.81	0.74	0.13	0.57	1.58	0.10	0.44	3.54	0.10	0.52	4.80	0.14
15	0.42	0.74	0.16	0.43	1.58	0.19	0.47	3.54	0.10	0.54	4.80	0.08
16	0.23	0.74	0.14	0.21	1.58	0.17	0.23	3.54	0.17	0.37	4.80	0.16
17	0.13	0.74	0.15	0.19	1.58	0.14	0.15	3.54	0.12	0.23	4.80	0.24
18	0.33	0.74	0.13	0.17	1.58	0.13	0.14	3.54	0.09	0.13	4.80	0.11
19	0.31	0.74	0.11	0.24	1.58	0.15	0.13	3.54	0.09	0.09	4.80	0.05
20	0.17	0.74	0.08	0.16	1.58	0.14	0.09	3.54	0.07	0.07	4.80	0.03

Note. $\mu = 100$, $\sigma = 10$, $b = 12$, $K = 800$, $h = 1$ (WW: Wagner Whitin, SM: Silver Meal, EIV: Ending Inventory Valuation).

than in the stationary case. As before, the SM heuristic generally performs worse when variability increases.

Correlated Demand

We generated a correlated demand pattern using a standard Markov demand process. We assume the existence of three demand states: high, medium, and low, with corresponding means $\mu_H = 140$, $\mu_M = 100$, and $\mu_L = 60$. We assume that demand in each state follows a normal distribution, and that the standard deviation of demand is the same for all states. As before, we assume that $K = 800$ and $h = 1$. In each period t, $1 \le t \le 300$, we assume that the state is observable, say i, and that the state in the next period will be j with probability $P_{i,j}$, where the transition probabilities are given in the following chart:

		State in period $t+1$		
		L	M	H
State in period t	L	0.70	0.25	0.05
	M	0.15	0.70	0.15
	H	0.05	0.25	0.70

The long-run average demand D is the expected demand in steady state, i.e., $D = \pi_H\mu_H + \pi_M\mu_M + \pi_L\mu_L$, where the πs denote the steady-state probabilities for each state. If the ending inventory typically covers several periods, then the long-run average demand may be an appropriate choice for ending-inventory calculation. In contrast, if the ending inventory typically covers only one period of demand, a more myopic demand estimate may be preferable. Specifically, for a model horizon of length T that starts in period t, the expected demand in period $t+T$ can be used for ending-inventory calculation. Assuming that the state in period $t+T-1$ is i, the expected demand in period $t+T$ is $D_{t+T} = P_{i,H}\mu_H + P_{i,M}\mu_M + P_{i,L}\mu_L$. Intuitively, the number of periods that the ending inventory will typically cover depends on the length of the model horizon relative to the number of periods covered by the average order (which depends on demand and setup and holding costs). For example, if the average order covers six periods and the model horizon is nine periods, we would expect ending inventory to cover about three periods of demand.

To compare the solution procedures in the case of correlated demand, we set $\mu_H = 140$, $\mu_M = 100$, and $\mu_L = 60$, and used four values for the standard deviation, $\sigma = 1, 5, 10$, and 20. As before, we set $K = 800$

Table 5a Percentage Deviation from Optimality For Correlated Demand, when the Myopic Demand Forecast Is Used for Ending Inventory Calculation

Model	$\sigma = 0$			$\sigma = 5$			$\sigma = 10$			$\sigma = 15$		
Horizon	WW	SM	EIV	WW	SM	EIV	WW	SM	EIV	WW	SM	EIV
2	32.10	32.10	14.21	32.37	32.37	16.36	32.89	32.89	16.14	33.64	33.64	16.79
3	7.65	7.65	11.02	7.82	7.82	15.85	8.20	8.20	16.90	8.76	8.76	17.52
4	2.85	1.97	6.06	2.99	2.04	6.41	3.32	2.44	6.07	3.83	2.77	5.97
5	5.53	0.71	1.03	4.95	0.82	1.18	5.72	1.25	1.47	5.56	1.61	1.75
6	3.80	0.75	1.16	4.20	0.88	1.31	4.48	0.99	1.34	4.72	1.14	1.38
7	3.10	0.75	1.05	2.48	0.95	0.96	2.54	1.04	0.80	2.47	1.16	0.91
8	1.48	0.75	0.51	1.57	0.95	0.52	1.35	1.04	0.43	1.66	1.16	0.49
9	1.35	0.75	0.50	1.31	0.95	0.54	1.28	1.04	0.55	1.34	1.16	0.42
10	1.23	0.75	0.32	1.02	0.95	0.42	1.07	1.04	0.30	1.04	1.16	0.29
11	0.94	0.75	0.21	0.82	0.95	0.22	0.70	1.04	0.26	0.62	1.16	0.27
12	0.78	0.75	0.22	0.69	0.95	0.26	0.62	1.04	0.22	0.63	1.16	0.16
13	0.44	0.75	0.16	0.58	0.95	0.17	0.48	1.04	0.16	0.42	1.16	0.14
14	0.41	0.75	0.15	0.31	0.95	0.13	0.29	1.04	0.12	0.22	1.16	0.07
15	0.31	0.75	0.10	0.23	0.95	0.10	0.27	1.04	0.08	0.29	1.16	0.09
16	0.29	0.75	0.09	0.25	0.95	0.08	0.20	1.04	0.08	0.20	1.16	0.05
17	0.16	0.75	0.07	0.25	0.95	0.06	0.15	1.04	0.06	0.11	1.16	0.04
18	0.19	0.75	0.07	0.15	0.95	0.06	0.11	1.04	0.04	0.16	1.16	0.03
19	0.18	0.75	0.07	0.13	0.95	0.04	0.10	1.04	0.02	0.11	1.16	0.02
20	0.11	0.75	0.04	0.09	0.95	0.03	0.09	1.04	0.03	0.12	1.16	0.03

Note. $\mu_H = 140$, $\mu_M = 100$, $\mu_L = 60$.

Table 5b Percentage Deviation from Optimality For Correlated Demand, when the Long-Term Average Demand Is Used for Ending Inventory Calculation

Model	$\sigma = 0$			$\sigma = 5$			$\sigma = 10$			$\sigma = 15$		
Horizon	WW	SM	EIV	WW	SM	EIV	WW	SM	EIV	WW	SM	EIV
2	31.81	31.81	11.63	31.96	31.96	16.39	32.33	32.33	16.35	32.95	32.95	17.05
3	7.70	7.70	13.11	7.83	7.83	19.36	8.14	8.14	19.67	8.65	8.65	20.08
4	2.77	2.22	2.22	2.92	2.34	2.34	3.26	2.70	2.65	3.78	3.03	3.02
5	5.80	1.23	1.18	5.12	1.11	1.42	5.60	1.33	1.36	5.84	1.86	1.50
6	3.86	0.89	0.76	4.01	1.03	0.78	4.06	1.21	1.05	4.37	1.32	1.17
7	3.10	0.89	0.60	2.51	1.06	0.85	2.62	1.21	0.81	2.45	1.29	0.65
8	1.59	0.89	0.56	1.64	1.06	0.47	1.45	1.21	0.41	1.66	1.29	0.50
9	1.48	0.89	0.47	1.41	1.06	0.37	1.28	1.21	0.34	1.15	1.29	0.27
10	1.15	0.89	0.27	1.00	1.06	0.29	0.96	1.21	0.21	1.11	1.29	0.23
11	0.98	0.89	0.21	0.71	1.06	0.18	0.75	1.21	0.23	0.68	1.29	0.18
12	0.62	0.89	0.19	0.70	1.06	0.19	0.62	1.21	0.18	0.50	1.29	0.21
13	0.48	0.89	0.15	0.55	1.06	0.15	0.41	1.21	0.13	0.39	1.29	0.12
14	0.42	0.89	0.13	0.37	1.06	0.08	0.26	1.21	0.07	0.23	1.29	0.05
15	0.30	0.89	0.10	0.28	1.07	0.09	0.25	1.22	0.09	0.21	1.30	0.09
16	0.19	0.89	0.08	0.24	1.07	0.09	0.20	1.22	0.09	0.25	1.30	0.05
17	0.20	0.89	0.05	0.20	1.07	0.06	0.19	1.22	0.05	0.15	1.30	0.03
18	0.17	0.89	0.05	0.12	1.07	0.05	0.10	1.22	0.05	0.08	1.30	0.03
19	0.14	0.89	0.03	0.13	1.07	0.05	0.08	1.22	0.04	0.07	1.30	0.02
20	0.04	0.89	0.02	0.09	1.07	0.05	0.10	1.22	0.04	0.10	1.30	0.03

Note. $\mu_H = 140, \mu_M = 100, \mu_L = 60$.

and $h = 1$. Average demand was calculated using both methods described above. For each set of input parameters, we ran eight demand simulations. The results are contained in Tables 5a and 5b. We find that the EIV significantly outperforms the WW and SM heuristics for all values of standard deviation for both methods of demand estimation used to calculate ending-inventory levels.

5. Concluding Remarks

In this paper, we show that end-effects in lot sizing can be mitigated by valuing ending inventory using a function that explicitly considers the tradeoff between holding inventory and incurring future setups. We developed this idea in the classical EOQ setting and demonstrated how it can be used in the context of the dynamic lot-sizing model. We empirically tested our method against the Wagner-Whitin algorithm and the well-regarded Silver Meal heuristic over a variety of demand patterns within a rolling horizon framework. For modest to long horizons, our method significantly outperforms WW and SM, except in the cases of extreme seasonality, steeply increasing demand, or steeply decreasing demand. In those cases, the differences were small. Our method does not perform as well for very short horizons that require only one setup. Intuitively, if the prescribed ending inventory is not enough to cover the first period after the horizon, there is no saving in setup cost to offset the additional carrying cost.

Our approach to ending-inventory valuation requires estimation of average demand beyond the model horizon. In practice, this estimation process could be prone to error. Research that tests the robustness of our approach with respect to estimation error for different types of demand distributions would enhance the practical value of our approach.

The usefulness of the concept of ending-inventory valuation to mitigate end-effects is not limited to the DLSP. It may be possible to use this concept for other variants of the uncapacitated single-item lot-sizing problem and for the more general multi-item capacitated lot-sizing problem. The latter problem has received much attention in the literature (e.g., Bitran and Matsuo 1986, Maes and Van Wassenhove

1986). In the multi-item capacitated lot-sizing problem, item production quantities in each period are linked through a shared capacity constraint. Trigiero et al. (1989) have provided a solution procedure for this problem that relaxes the linking capacity constraints. If the capacity constraints are dualized with appropriate shadow prices attached to them, the relaxed problem decomposes into a set of dynamic lot-sizing problems with time varying unit costs. In each period, the unit cost reflects the holding cost as well as the opportunity cost of capacity, which would be high in a period of peak demand. The inclusion of an ending-inventory valuation term that is modified for time varying unit costs in the objective function of each such subproblem may mitigate end-effects in multi-item capacitated lot sizing. However, with capacity constraints we may want to keep ending inventory to relieve a capacity shortfall after the short model horizon T. It may be possible to extend the concept of ending-inventory valuation to achieve this end.

Acknowledgments
We are grateful to three anonymous referees for their suggestions and insights, to Sanjay Jain for comments on an earlier version of this paper, and to Nils Rudi for help with coding the simulation tests.

References

Aggarwal, A., J. K. Park. 1993. Improved algorithms for economic lot size problems. *Oper. Res.* **41** 549–571.

Baker, K. R., P. Dixon, M. J. Magazine, E. A. Silver. 1978. An algorithm for the dynamic lot size problem with time-varying production capacity constraints. *Management Sci.* **24** 1710–1720.

Bitran, G., H. Matsuo. 1986. The multi item capacitated lot size problem: Error bounds of Manne's formulation. *Management Sci.* **32** 350–359.

Blackburn, J. D., R. A. Millen. 1980. Heuristic lot sizing performance in a rolling-schedule environment. *Decision Sci.* **11** 691–701.

Carr, C.R. and C.W. Howe. 1962. Optimal service policies and finite time horizons. *Management Sci.* **9**(1) 126–140.

Chand, S., S. Sethi, G. Sorger. 1992. Forecast horizons in the discounted dynamic lot size model. *Management Sci.* **38** 1034–1048.

Chen, H.-D., D. W. Hearn, C.-Y. Lee. 1995. Minimizing the error bound for the dynamic lot size model. *Oper. Res. Letters* **17** 57–68.

Eppen, G. D., F. J. Gould, B. P. Pashigian. 1969. Extensions of planning horizon theorem in the dynamic lot size model. *Management Sci.* **15** 268–277.

Federgruen, A., M. Tzur. 1991. A simple forward algorithm to solve general dynamic lot sizing models in $O(n \log n)$ or $O(n)$ time. *Management Sci.* **37** 909–925.

———, ———. 1994. Minimal forecast horizons and a new planning procedure for the general dynamic lot sizing model: Nervousness revisited. *Oper. Res.* **42** 456–468.

Grinold, R. C. 1983. Model building techniques for the correction of end effects in multistage convex programs. *Oper. Res.* **31** 407–431.

Lee, C. Y., E. V. Denardo. 1986. Rolling planning horizons: Error bounds for the dynamic lot size model. *Math. Oper. Res.* **11** 423–432.

Lundin, R. A., T. E. Morton. 1975. Planning horizons for the dynamic lot size model: Zabel vs. protective procedures and computational results. *Oper. Res.* **23** 711–734.

Maes, J., L. N. Van Wassenhove. 1986. Multi item single level capacitated dynamic lot sizing heuristics: A computational comparison. *IIE Trans.* 124–129.

Silver, E. A., H. C. Meal. 1973. A heuristic selecting lot size requirements for the case of a deterministic time-varying demand rate and discrete opportunities for replenishment. *Production and Inventory Management* **14** 64–77.

Trigeiro, W. W., L. J. Thomas, J. O. McClain. 1989. Capacitated lot sizing with setup times. *Management Sci.* **35**(3) 353–366.

Vachani, R. 1992. Performance of heuristics for the uncapacitated lot size problem. *Naval Res. Logist.* **39** 801–813.

Wagelmans, A., S. Van Hoesel, A. Kolen. 1992. Economic lot sizing: An $O(n \log n)$ algorithm that runs in linear time in the Wagner Whitin case. *Oper. Res.* **40** 8145–8156.

Wagner, H. M., T. M. Whitin. 1958. Dynamic version of the economic lot size model. *Management Sci.* **5** 89–96.

Zabel, E. 1964. Some generalization of inventory planning horizon theorem. *Management Sci.* **10** 465–471.

Accepted by Luk Van Wassenhove; received May 2, 1995. This paper was with the authors 44 months for 2 revisions.