MINIMAL FORECAST HORIZONS AND A NEW PLANNING PROCEDURE FOR THE GENERAL DYNAMIC LOT SIZING MODEL: NERVOUSNESS REVISITED

AWI FEDERGRUEN

Columbia University, New York, New York

MICHAL TZUR

University of Pennsylvania, Philadelphia, Pennsylvania

(Received January 1991; revisions received July 1991, March 1992; accepted October 1992)

We show for the general dynamic lot sizing model how minimal forecast horizons may be detected by a slight adaptation of an earlier $O(n \log n)$ or $O(n)$ forward solution method for the model. A detailed numerical study indicates that minimal forecast horizons tend to be small, that is, include a small number of orders. We describe a new planning approach to ensure stability of the lot sizing decisions over an initial interval of time or *stability horizon* in those (relatively rare) cases where no planning horizon is detected or where the stability horizon extends beyond the planning horizon. To this end, we develop a heuristic, but *full horizon-based* adaptation of the optimal lot sizing schedule, designed to minimize an upper bound for the worst-case optimality gap under the desired stability conditions. We also show how the basic horizon length n may be chosen to guarantee any prespecified positive optimality gap.

The dynamic lot sizing model is one of the most commonly used (single-item) production/inventory planning models: For a given horizon of n periods, one needs to find a schedule of order quantities or lot sizes, to satisfy a given demand in each of the periods while minimizing the sum of fixed-plus-variable order and linear inventory carrying costs. All cost and demand parameters may be time dependent. The dynamic lot sizing model arises as a repeatedly solvable subproblem in, for example, material requirement planning (MRP) systems, hierarchical planning problems, and multi-item capacitated lot sizing problems; see Federgruen and Tzur (1991) for a more detailed discussion.

Wagner and Whitin (1958) showed that an *exact* solution of the dynamic lot sizing problem can be found with a relatively simple $O(n^2)$ shortest path algorithm. Subsequently, a large volume of heuristics was developed, see, for example, the least unit cost heuristic, the part-period balancing heuristic, the economic order quantity heuristic, as well as Silver and Meal (1973), Peterson and Silver (1979), Axsäter (1982, 1985), Bitran, Magnanti and Yanasse (1984), and the survey of Baker (1989). These heuristics were designed for special cases of the lot sizing model, for example, with constant fixed and per unit order costs. Moreover, the decision of whether to order in a given period is, in virtually all of these heuristics, based on *prior* demand and cost data only, thus ignoring information about later periods entirely. Therefore, we refer to these heuristics as *history-based* procedures as opposed to *full horizon-based* lot sizing decisions which arise in exact solution methods.

The need for such heuristics has been motivated on the grounds that:

i. the $O(n^2)$ complexity of the Wagner-Whitin algorithm was considered prohibitive for many of the above-mentioned settings where the model needs to be solved repeatedly for a larger number of items and/or different combinations of parameter values;
ii. as is the case with most dynamic planning problems, the dynamic lot sizing model is usually solved on a *rolling horizon* basis, whereby as each period passes, it is eliminated and replaced by a new period, appended at the end of the horizon. Many have asserted that a new optimal schedule obtained by exact optimization over the most recent horizon (with possibly updated parameters), may differ markedly from the previous schedule upon which some plans may already have been

Subject classifications: Inventory/production: policies, planning horizons, and sensitivity analysis.
Area of review: MANUFACTURING, OPERATIONS AND SCHEDULING.

0030-364X/94/4203-0456 $01.25

based. This phenomenon is referred to as *nervousness* (see Steele 1973).

It is argued that nervousness is best eliminated by adopting one of the history-based heuristics. In this paper, we provide a new perspective on the use of history-versus full horizon-based methods, and describe how the latter can be designed efficiently to deal with the above-mentioned problems.

As far as the complexity issue is concerned, new exact algorithms have been developed by Federgruen and Tzur (1991), Wagelmans, Van Hoesel and Kolen (1992) and Aggarwal and Park (1993) whose complexity is $O(n \log n)$ for the general model and $O(n)$ for all cases for which the history-dependent heuristics were designed. In other words, there now exist exact (full horizon-based) solution methods whose complexity is of the exact same order of magnitude as that of these heuristics.

The potential for *nervous* lot sizing schedules arises because of *two* distinct causes: the *revised data effect* and the *truncated horizon effect*. The former arises in rolling horizon procedures as one proceeds from one solution of the lot sizing model to the next, because cost as well as demand parameters are modified on the basis of new and more accurate data. The second effect consists of the extension of the horizon, potentially resulting in a change of many or all lot sizes, including the ones that pertain to initial periods.

As far as the revised data effect is concerned, history-based heuristics are no more immune to the resulting potential for nervousness than their full horizon-based counterparts. The former appear to have the advantage that the impact of a revised demand estimate for a given period is restricted to that and future periods only. The same partial stability may, however, be achieved in any full-horizon based method by fixing the decisions in all periods prior to those with revised demands. We now discuss why the latter approach results in more cost-effective schedules while ensuring the same (partial) stability.

Several authors have addressed the *truncated horizon effect*, by describing procedures for the detection of *planning* and *forecast horizons*. A period L is called a *forecast horizon* if the decisions for the first ℓ periods in the optimal solution of a problem with horizon length L are not affected by the model parameters for periods beyond the horizon L. Such a period ℓ is called a *planning horizon*. We refer to n as the *study* horizon. (Some authors use the term decision horizon for planning horizon and planning horizon for study horizon.) Clearly, all periods in a planning horizon are immune to the truncated horizon effect.

Existing procedures for the identification of a planning horizon have all been confined to special cases of the model, except for Bensoussan, Proth and Queyranne (1991) in which a procedure is described for the identification of a maximal planning horizon (if any) for any *given* study horizon. Their algorithm handles general, piecewise linear concave order and holding cost functions, a generalization of the basic model. The complexity of their algorithm is $O(n^3)$ when applied to the basic model (and of higher complexity under general, piecewise linear functions).

We show how *minimal* forecast horizons may be detected in the *general* model by a slight adaptation of the $O(n \log n)$ forward algorithm in Federgruen and Tzur. (This appears to be a unique advantage of this algorithm compared with the other $O(n \log n)$ alternatives.) Moreover, a detailed numerical study indicates that forecast horizons tend to be small, i.e., include a few orders, in particular, under considerable variability in the cost or demand data.

The final question relates to ensuring stability of lot sizing decisions over an initial interval of time or *stability horizon* in those (relatively rare) cases where no planning horizon is detected or when this stability horizon extends beyond the planning horizon. Here we develop a heuristic, but *full horizon-based* adaptation of the optimal lot sizing schedule, designed to minimize an upper bound for the worst-case optimality gap under the desired stability conditions. We also show how the basic horizon length n may be chosen to guarantee any prespecified, positive optimality gap.

In many practical single or multi-item lot sizing problems, it is necessary to address one or more capacity constraints. Even for the single-item capacitated problem, no exact forecast horizon procedures are known. Indeed, the problem is NP-complete (see Florian, Lenstra and Rinnooy Kan 1980). This is why capacitated models are transformed into uncapacitated ones in practice. For example, even the more advanced material requirement planning systems (MRPII) start determining system-wide order releases without considering capacity constraints, i.e., on the basis of the uncapacitated single-item model, treated here. Only at the end is an attempt made to eliminate capacity conflicts by heuristic adaptations of the schedules. It is, therefore, important to demonstrate (whenever possible) that at least the basic schedules are invariant with respect to unknown future demand or cost parameters. Moreover, even for the single-item capacitated model, many heuristics are based on heuristic adaptations of an optimal schedule for the uncapacitated version (see, for example, Dixon and Silver 1981, Karni 1981, and Nahmias 1989, subsection 6.6).

In other settings the uncapacitated model arises as a subproblem in a more complex (for example, hierarchical) planning model which is solved on a rolling horizon basis via Lagrangian relaxation. The cost parameters in such applications depend on the Lagrange multipliers. The bound maximizing values of these multipliers (even those which pertain to an initial interval of periods) vary in general with the chosen horizon. Forecast horizons found for the single-item subproblems and a given vector of Lagrange multipliers imply that the decisions which pertain to the associated planning horizons are optimal for *all* parameter values *beyond* the forecast horizons, as well as sufficiently small variations of the Lagrange multipliers relevant to periods *preceding* these forecast horizons. (Exact ranges for the permissible variation of individual Lagrange multipliers are easy to compute, using the results in Van Hoesel and Wagelmans 1993.) The optimal Lagrange multipliers for initial periods are likely to be insensitive with respect to the chosen study horizon, as long as the overall capacity utilization is not close to one. The planning horizons identified for the uncapacitated models thus imply that decisions for periods up to these horizons remain optimal, as long as the optimal Lagrange multipliers for periods preceding the forecast horizon lie within the above-mentioned ranges, or, the search procedure for Lagrange multipliers is restricted to these ranges.

Algorithms to detect planning and forecast horizons are due to Wagner and Whitin (1958), Zabel (1964), Lundin and Morton (1975), Chand (1982) and Bensoussan, Crouhy and Proth (1983). These algorithms apply to special cases where no speculative motives exist, i.e., where the per unit order cost increases by no more than the cost of carrying a unit in stock in each interval of time. Chand, Sethi and Proth (1990) and Chand, Sethi and Sorger (1992) address the special case where *all* parameters are constant, establishing existence as well as upper bounds for a forecast horizon in the undiscounted and discounted model, respectively. Eppen, Gould and Pashigian (1969) develop planning horizon procedures for the *general* model. These procedures may all fail to detect forecast horizons, i.e., they apply test conditions which are sufficient but not necessary. Only Chand and Morton (1986) describe a procedure for detecting a *minimal* forecast horizon for models without speculative motives. (In this case, the optimal, last setup period is nondecreasing in the horizon length, as observed in Wagner and Whitin, thus markedly facilitating the detection of forecast horizons.)

A crucial element in our forecast horizon procedure is the construction of a *minimal* list of periods which are optimal, *last* order periods for some conceivable future horizon. Lundin and Morton first introduced a similar concept of a *regeneration set*, defined as a set containing optimal regeneration points (among possibly others), i.e., periods followed by an order, in some conceivable future horizon. Their planning horizon theorems contain sufficient conditions for detecting a planning horizon, stated in terms of these regeneration sets; Lundin and Morton's work also appears to be the first to address optimality gaps associated with the first lot sizing decision in settings where no forecast horizon has been detected.

As mentioned, Bensoussan, Proth and Queyranne describe a procedure for **MP**, the identification of a *maximal planning* horizon (if any) for a given study horizon n. The procedures in this paper identify *minimal forecast* (**MF**) horizons for a given stability horizon $s \geq 1$. All other existing previous work addresses the same problem, albeit for the case $s = 1$ only. **MP** and **MF** may be viewed as *dual*. Indeed, a procedure for **MP** may be transformed into a procedure for **MF** and vice versa, via a binary search with respect to the study horizon n or the stability horizon s, respectively. Lee and Denardo (1986) determine, for models *without* speculative motives, upper bounds for the error (i.e., the extra costs) which may be incurred by fixing the first ℓ periods' lot sizes on the basis of a lot sizing model of given length n, rather than some larger length $T > n$. The results in this paper extend theirs to the *general* lot sizing model.

The remainder of this paper is organized as follows: In Section 1 we derive our procedures for the detection of minimal forecast horizons and report the above-mentioned numerical study. In Section 2 we derive bounds for the optimality gap associated with any given first period decision; this section also describes our proposed full horizon-based planning procedure.

1. PROCEDURES FOR FINDING MINIMAL FORECAST AND PLANNING HORIZONS

In this section we describe procedures for finding minimal forecast horizons (and associated planning horizons) for the *general* dynamic lot sizing model. The derivation of these procedures is based on an $O(n \log n)$ algorithm *recently* developed by Federgruen and Tzur to solve the dynamic lot sizing problem.

The model is described by the parameters:

d_i = the demand in period i ($i = 1, \ldots, n$);
K_i = the setup cost in period i ($i = 1, \ldots, n$);
c_i = the variable per unit order cost in period i ($i = 1, \ldots, n$);
h_i = the cost of carrying a unit of inventory at the end of period i ($i = 1, \ldots, n$);

We assume, without loss of generality, that the starting inventory in period 1 and the ending inventory in period n equal *zero*.

We use the following auxiliary notation:

$D(i) = \sum_{k=1}^{i} d_k$ represents the cumulative demand in periods $1, \ldots, i$.

For all $i < j$, let

$d_{ij} = d_i + \cdots + d_{j-1}$
$c_{ij} = c_i + h_i + \cdots + h_{j-1}$
(the variable cost of ordering a unit in period i and carrying it till period j);
$C(i) = c_{in}$.

(Thus, with some algebra one verifies that for all $i < j$ the following equivalency is satisfied: $c_{ij} \leq c_j$ if and only if $C(i) \leq C(j)$.)

In subsection 1.1 we summarize the properties of this algorithm that are essential for the derivation of the forecast horizon procedures. In subsection 1.2 we develop the procedure under the assumption that the optimal last order period for each horizon $j = 1, \ldots, n$ is *unique*. We then relax this assumption in subsection 1.3 and extend the procedure to the general case, allowing for multiple, optimal last production periods. In subsection 1.4 we consider models in which at the beginning of period 1 all cost parameter trajectories are assumed to be known ad infinitum. Since we assume that no parameter values beyond the forecast horizon are known in the general case, a minor modification to the general procedure is required. Finally, in subsection 1.5 we extend the minimal forecast horizon procedures to ensure that the associated planning horizon is at least as large as a prespecified so-called *stability horizon*. Our computational results are presented in subsection 1.6.

1.1. Summary of an $O(n \log n)$ Algorithm to Solve the Dynamic Lot Sizing Problem

It is well known since Wagner and Whitin that optimal policies exist under which orders are placed if and only if inventory equals zero, and such *zero-inventory ordering policies* are completely determined by the specification of the last order period $\ell(j)$ preceding any given horizon j ($j = 1, \ldots, n$). Thus, we only consider solutions that satisfy the zero inventory property. Let $F(j)$ = the minimum cost in the first j periods ($j = 1, \ldots, n$), and $F(\ell, j)$ = the minimum cost in the first j periods if the final setup is performed in period $\ell \leq j$ ($j = 1, \ldots, n$).

The above-mentioned algorithm is a forward procedure with n iterations. In the jth iteration we treat all future parameter values pertaining to periods ($j + 1, j + 2, \ldots, n$) as unknown so that, in particular, all future cumulative demands may adopt any potential value $D \geq D(j)$. The algorithm constructs in the jth iteration a list $\Omega(j)$ containing all periods that, among the first j periods, are optimal last order periods for some potential horizon $t \geq j$ with cumulative demand $D \geq D(j)$. More specifically:

Definition. Let $\Omega(j) = \{i_1, i_2, \ldots, i_m\} = \{1 \leq \ell \leq j$: There exists a horizon $t \geq j$ with a potential cumulative demand $D \geq D(j)$ for which ℓ is the *lowest* index with $F(\ell, t) = \min_{1 \leq i \leq j} F(i, t)\}$. Associated with the list $\Omega(j)$ are values $D(j) = g(1) < \cdots < g(m)$ with the following property.

Property. Let $\Omega(j) = \{i_1, i_2, \ldots, i_m\}$. For all $k = 1, \ldots, m$, i_k is an optimal last order period among the periods $1, \ldots, j$ for a horizon $t \geq j$ *if and only if* D, the cumulative demand over that horizon satisfies $g(k) \leq D \leq g(k + 1)$ (with the convention $g(m + 1) = \infty$). (See Lemma 2 and the discussion in Section 2 in Federgruen and Tzur.)

The lists $\Omega(j)$ ($j = 1, \ldots, n$) are recursively updated and at the end of the jth iteration $\ell(1), \ldots, \ell(j)$ are available. The following inclusion holds:

$$\Omega(j + 1) \subseteq [\Omega(j) \cup \{j + 1\}] \quad j = 1, \ldots, n - 1. \quad (1)$$

1.2. Minimal Forecast Horizon Procedure Under Unique Optimal Solutions

In this subsection we assume that $\ell(j)$ is unique for all $j = 1, \ldots, n$. This assumption implies that a *unique schedule* exists for all horizons $j = 1, \ldots, n$. We derive an efficient procedure for the identification of the shortest forecast horizon and its associated planning horizon. Let

$q(j)$ = the *first* order period after period 1 in any given schedule which is optimal among all schedules that place an order in period j ($j = 2, \ldots, n$);
$q(1) = 1$.

(Note that in every optimal schedule which places an order in period j, the ending inventory in period $j - 1$ equals zero; thus, all such schedules prescribe the same decisions in periods $1, \ldots, j - 1$.)

We first derive the necessary and sufficient condition for a period j to be a forecast horizon. As pointed out in the Introduction, this generalizes Chand and Morton's necessary and sufficient condition for models without speculative motives for carrying inventories.

Theorem 1. *Fix* $j = 1, \ldots, n$. *Let* $\Omega(j) = \{i_1, i_2, \ldots, i_m\}$. *Period* j *is a forecast horizon if*

$$1 < q^* \equiv q(i_1) = q(i_2) = \cdots = q(i_m). \quad (2)$$

If (2) *holds, then* $q^* - 1$ *is a planning horizon.*

We first need the following lemma.

Lemma 1. *Fix* $j \geq 1$. *In the optimal schedule for any horizon* $t \geq j$, *an order is placed in some period* $i \in \Omega(j)$.

Proof. By induction with respect to t. By the definition of $\Omega(j)$ the lemma clearly holds for $t = j$; assume that it holds for *all* horizons $j, j + 1, \ldots, t$. Consider now the horizon of length $t + 1$. If $\ell(t + 1) \leq j$, then $\ell(t + 1) \in \Omega(j)$, by the definition of $\Omega(j)$, and the lemma holds. If $j < \ell(t + 1) \leq t + 1$, it follows from the zero-inventory ordering property that the first part of the optimal schedule for the horizon of length $t + 1$ is also optimal for the horizon of length $[\ell(t + 1) - 1] \leq t$. The lemma now follows from the induction assumption.

Proof of Theorem 1. Assume first that (2) holds and consider an arbitrary horizon $t \geq j$. By Lemma 1, the optimal schedule for this horizon prescribes an order in some period $i \in \Omega(j)$ and hence in $q(i) = q^*$. By the zero-inventory ordering property, it is therefore optimal to order $D(q^* - 1)$ units in period 1, and nothing in periods $2, \ldots, q^* - 1$, regardless of cost and demand parameters pertaining to periods after period j. Therefore, j is a forecast horizon and $q^* - 1$ the associated planning horizon.

Conversely, assume that (2) fails to hold but that period j is a forecast horizon nevertheless. This implies, in particular, that the optimal order quantity in the first period is independent of any parameter value that pertains to periods after period j. Thus, in view of the zero-inventory ordering property, assume that it is optimal to order in period 1 the cumulative demand up to (but not including) some period $r \geq 1$ for *any* horizon $t \geq j$. This implies that period r is the first order period after period 1 in the corresponding optimal schedule. Since (2) fails to hold, there exists an element $i_k \in \Omega(j)$ with $q(i_k) \neq r$. Assume now that the *unknown* cumulative demand $D(j + 1)$ satisfies $g(k) < D(j + 1) < g(k + 1)$ and either K_{j+1} or c_{j+1} is sufficiently large to preclude ordering in period $j + 1$. By the Property, $\ell(j + 1) = i_k$ and it is *uniquely* optimal to order in period 1 the cumulative demand up to (but not including) period $q(i_k) \neq r$, which is a contradiction.

Theorem 1 suggests an extremely simple procedure for the detection of a minimal forecast horizon and its associated planning horizon. Note, once again by the zero-inventory property, that

$$q(j) = \begin{cases} j & \text{if } \ell(j - 1) = 1 \\ & (j = 2, \ldots, n). \\ q(\ell(j - 1)) & \text{otherwise} \end{cases} \quad (3)$$

The above-mentioned forward solution method of the dynamic lot sizing problem is thus easy to adapt to allow for the detection of a minimal forecast horizon. Having computed and stored the values $q(2), \ldots, q(j - 1)$ in the first $j - 1$ iterations, $q(j)$ is easy to compute via (3). Thus, to verify whether period j is a forecast horizon, it suffices to conclude the jth iteration with a test of whether (2) holds or not. This can be done efficiently by keeping track of the multiplicity of each distinct q-value among all q-values associated with elements of Ω.

Thus, using an additional n-array COUNT, let

$$\text{COUNT}[r] = |\{\ell \in \Omega | q(\ell) = r\}| \quad (r = 1, \ldots, n).$$

Note that this array needs to be updated only when an element is added to or deleted from Ω, in which case, *one* of its entries is increased or decreased by one, respectively. At the end of the jth iteration of the algorithm, condition (2) may thus be checked by testing whether for example,

$$\text{COUNT}[q(i_1)] = |\Omega(j)|. \quad (4)$$

The work required to update the COUNT-array is constant per single deletion or addition of an element of Ω. There are, at most, $2n$ such deletions and additions (see Federgruen and Tzur). The total amount of *additional* work required when appending the minimal forecast horizon procedure to the basic algorithm, is thus $O(n)$ (including the test (4) at the end of each iteration). In other words, the *integrated* algorithm continues to have complexity $O(n \log n)$ and $O(n)$ in the special cases identified in Federgruen and Tzur.

1.3. Minimal Forecast Horizon Procedure Under Multiple Optimal Solutions

In this subsection, we consider the general model in which several optimal last order periods and, hence, several optimal schedules may exist for some horizons j ($j = 1, \ldots, n$). We thus define:

$$L(j) = \{\ell: F(\ell, j) = F(j)\} \quad (j = 1, \ldots, n)$$

i.e., $L(j)$ is the set of *all* optimal last order periods in the problem with horizon length j. Also, let

$Q(j) = \{\ell$: ℓ is the *first* order period, after period 1, in *some* schedule which is optimal among all schedules that place an order in period $j\}$ ($j = 2, \ldots, n$),

$Q(1) = \{1\}$

One easily verifies, in direct analogy to (3), that

$$Q(j) = \begin{cases} \bigcup_{r \in L(j-1)} Q(r) & \text{if } 1 \notin L(j-1) \\ & j = 2, \ldots, n. \\ \bigcup_{r \in L(j-1) \setminus \{1\}} Q(r) \cup \{j\} & \text{otherwise} \end{cases} \quad (5)$$

Computation of the sets $L(j)$ and $Q(j)$ ($j = 2, \ldots, n$) requires a slight modification of the algorithm in Federgruen and Tzur. The algorithm there, designed to identify *one* optimal solution only, eliminates elements from the candidate list Ω, whenever *at least as good* an alternative last order period prevails for every conceivable future horizon. To identify the sets $L(j)$, $j = 1, \ldots, n$ one needs to maintain *every* element which is not strictly dominated for all future horizons. It follows from the algorithm's stated property (see subsection 1.1) that an element $i_k \in \Omega(j)$ is the *unique* optimal last order period for some potential future horizon if $g(k) < g(k + 1)$. If $g(k) = g(k + 1)$ however, i_k is a (nonunique) optimal last order period for a future horizon with cumulative demand $D = g(k) = g(k + 1)$ and it is strictly dominated by some other element in the list $\Omega(j)$ for all other values of D.

The algorithm in Federgruen and Tzur eliminates such nonessential periods from the list. For the purpose of identifying the complete sets $L(j)$ ($j = 1, \ldots, n$) it is, however, necessary to keep such periods in the list. Thus, with this modification, we construct a list $\bar{\Omega}(j) = \{\ell_1, \ldots, \ell_{\bar{m}}\}$ defined by: $\{1 \leq \ell \leq j$: There exists a horizon $t \geq j$ with a potential cumulative demand $D \geq D(j)$ for which $F(\ell, t) = \min_{1 \leq i \leq j} F(i, t)\}$ and with associated $\bar{g}$-values such that $\bar{g}(1) \leq \bar{g}(2) \leq \cdots \leq \bar{g}(\bar{m})$, as opposed to the list $\Omega(j)$ constructed in the original algorithm in which the associated g-values are strictly ordered. (Note that $\bar{\Omega}(j) \supseteq \Omega(j)$, $j = 1, \ldots, n$.) Thus, assume that the $\bar{g}$-values associated with a list $\bar{\Omega}(j) = \{\ell_1, \ldots, \ell_{\bar{m}}\}$ contain m *distinct* g-values $g(1) < \cdots < g(m)$ with $\bar{m} \geq m$. For $k = 1, \ldots, m$, let $i_k = \ell_{\max\{r: \bar{g}(r) = g(k)\}}$ denote the highest indexed element of $\Omega(j)$ with the kth smallest g-value (see Table I). (Note from the property that among all elements of $\bar{\Omega}(j)$ with $\bar{g}$-value equal to $\bar{g}(r)$ only i_k is a *unique*, optimal last order period for certain potential future horizons. In other words, the minimal list $\Omega(j) = \{i_1, \ldots, i_m\}$.)

With the modified algorithm, and in view of the property, we easily identify $L(j)$ at the end of the jth iteration as $L(j) = \{\ell_k \in \bar{\Omega}(j): \bar{g}(k) = \bar{g}(1) = D(j)\}$. Having computed $L(j)$, $Q(j)$ is identified via (5).

We now restate the necessary and sufficient condition for a period j to be a forecast horizon.

Theorem 2. *Fix $j = 1, \ldots, n$. Period j is a forecast horizon iff there exists a period*

$$1 < q^* \in Q(i_r) \quad \textit{for all } r = 1, \ldots, m. \quad (6)$$

If (6) holds, $q^* - 1$ is a planning horizon.

Proof. In the general case where multiple optimal schedules may exist Lemma 1 continues to hold for *some* optimal schedule. Fix $1 < q^* \in Q^* \equiv \cap\{Q(i_r)$: $r = 1, \ldots, m\}$. The proof of the sufficiency part of (6) is identical to that given for Theorem 1.

Conversely, assume that (6) fails to hold, but that period j is a forecast horizon nevertheless. Define period r as in the necessity part of the proof of Theorem 1. Since (6) fails to hold, there exists an element $i_k \in \Omega(j)$ such that $r \notin Q(i_k)$. Specify period ($j + 1$)'s parameters as in the proof of Theorem 1, so that in an optimal schedule for the corresponding planning problem with $j + 1$ periods, i_k is the *unique*, best last production period. In other words, by the definition of the set $Q(i_k)$, for the schedule to be optimal it is necessary to order in period 1 up to some period in $Q(i_k)$, but $r \notin Q(i_k)$, which is a contradiction.

Theorem 2 suggests a simple procedure for the detection of a minimal forecast horizon for the general case where multiple optimal solutions may exist. To

Table 1
$\bar{\Omega}(\cdot) \supseteq \Omega(\cdot)$

Index	g-Value
ℓ_1	$\bar{g}(1)$
ℓ_2	$\bar{g}(2) = \bar{g}(1)$
$\ell_3 = i_1$	$\bar{g}(3) = \bar{g}(1) = g(1)$
.	.
.	.
.	.
ℓ_r	$\bar{g}(r)$
$\ell_{r+1} = i_k$	$\bar{g}(r + 1) = \bar{g}(r) = g(k)$
.	.
.	.
$\ell_{\bar{m}-1}$	$\bar{g}(\bar{m} - 1)$
$\ell_{\bar{m}} = i_m$	$\bar{g}(\bar{m}) = \bar{g}(\bar{m} - 1) = g(m)$

verify whether period j is a forecast horizon it suffices to conclude the jth iteration of the (modified) algorithm with a test of whether (6) holds. As before, this can be done efficiently by keeping track of the multiplicity of each of the q-values that belong to one of the Q-sets associated with elements of Ω. The same COUNT-array may be used; its updates proceed as before and (6) may be verified at the end of the jth iteration by checking whether for some $q \in Q(i_1)$, COUNT$[q] = |\Omega(j)|$. Under the extremely innocuous assumption that the cardinality of the Q-sets is uniformly bounded, we conclude again that the additional work required when appending the minimal forecast horizon procedure to the basic algorithm, remains $O(n)$. The *integrated* algorithm thus has complexity $O(n \log n)$ in the *most general* case and $O(n)$ in the special cases identified in Federgruen and Tzur.

1.4. Minimal Forecast Horizon Procedures Under Known Cost Trajectories

In this subsection, we describe a modified procedure for the detection of *minimal* forecast horizons in models with known cost trajectories, i.e., where at the beginning of period 1 the parameters K_t, h_t, and c_t are assumed to be known for all $t \geq 1$. Special cases include models with constant cost parameters, as treated, for example, by Chand, Sethi and Sorger, and Chand, Sethi and Proth. Other cases include settings where the cost parameters follow a given trend and/or cyclical pattern. Our analysis indicates that to establish that a given period j is a forecast horizon, of all future cost parameters, only those pertaining to period $j + 1$ can be used. Chand and Morton also give special treatment to this case, for models without speculative motives for carrying inventories.

For the sake of notational simplicity, we restrict ourselves to the case where a unique optimal schedule exists for all horizons $j = 1, \ldots, n$. (The extension to the general case is straightforward along the lines of subsection 1.3.)

Note first that (2) continues to be sufficient. It may, however, fail to be necessary because under known cost trajectories some of the elements in the list $\Omega(j)$ can *never* arise as order periods in the optimal schedule of *any* future horizon $t \geq j$. In constructing $\Omega(j + 1)$ in the $(j + 1)$st iteration of our forward solution method, some elements of $\Omega(j)$ are eliminated because, after the addition of period $j + 1$, they are identified as being dominated as last order periods in *any* $(j + 1)$-period problem with potential cumulative demand $D \geq D(j)$. The elimination tests require knowledge of period $(j + 1)$st cost parameters but, by definition, *not* of the future demand value d_{j+1}. Under the assumption of known cost trajectories, these eliminations can thus be conducted on the basis of the information available up to period j. Let $\underline{\Omega}(j) \subseteq \Omega(j)$ denote the list obtained after the above deletions. (These deletions correspond with the execution of the procedures DELUP and DELDOWN in Step 1 of the algorithm in Federgruen and Tzur.) Using the Property, it is now easy to verify that *every* $i \in \underline{\Omega}(j)$ is the *unique* optimal last order period for the problem with horizon length $j + 1$ and *some* cumulative demand D. The following minimal forecast horizon test now follows directly from the proof of the necessity part of Theorem 1. The test generalizes that obtained in Chand and Morton for models without speculative motives for carrying inventories.

Theorem 3. *Assume that all cost trajectories c_i, K_i, h_i are known for all $i = 1, 2, \ldots$. Fix $j = 1, \ldots, n$. Let $\underline{\Omega}(j) = \{k_1, \ldots, k_{\underline{m}}\}$ (with $\underline{m} \leq m$). Period j is a forecast horizon iff*

$$1 < q^* \equiv q(k_1) = q(k_2) = \cdots = q(k_{\underline{m}}) \tag{7}$$

If (7) holds, period $q^* - 1$ is a planning horizon.

As before, the additional work required to perform test (7) at the end of each iteration remains of complexity $O(n)$. The integrated algorithm (which includes a procedure for the detection of a minimal forecast horizon) thus continues to require $O(n \log n)$ operations. In special cases, for example, those with *constant* cost parameters, the algorithm may be implemented in linear time (see Sections 3 and 4 in Federgruen and Tzur).

1.5. Minimal Forecast Horizon Procedures for a Given Stability Horizon

To our knowledge, all forecast horizon procedures, including the ones derived in subsections 1.2–1.4, accept any associated planning horizon. Such procedures are perfectly adequate if at the beginning of period 1 a commitment needs to be made with respect to the first period only. However, in many practical settings it is required or desired to obtain a planning horizon which is at least as large as a prespecified interval of time, which we refer to as the *stability horizon* (see the Introduction).

In this subsection we describe how our minimal forecast horizon procedure may be extended to ensure that a given stability horizon s is achieved. As in the previous subsection and for notational simplicity only, we restrict ourselves to the case where a unique optimal schedule exists for all horizons $j = 1, \ldots, n$.

(The extension to the general case is again straightforward along the lines of subsection 1.3.)

For a given stability horizon $s \geq 1$, define:

$$q_s(j) = \begin{cases} \text{the } \textit{first} \text{ order period after period } s \text{ in any schedule which is optimal among } \textit{all} \text{ schedules that place an order in period } j & j = s+1, \ldots, n \\ 1 & \text{otherwise.} \end{cases}$$

Note that $q_1(j) = q(j)$ for all $j = 1, \ldots, n$. In analogy with (3) one verifies that:

$$q_s(j) = \begin{cases} j & \text{if } \ell(j-1) \leq s \\ & (j = s+1, \ldots, n) \\ q(\ell(j-1)) & \text{otherwise.} \end{cases} \tag{8}$$

The following corollary describes how to detect minimal forecast horizons, guaranteeing a planning horizon of length s or larger. Its proof is immediate from Theorem 1 and the definition of $q_s(\cdot)$.

Corollary 1. *Fix $j = 1, \ldots, n$ and $s \leq j$. Let $\Omega(j) = \{i_1, \ldots, i_m\}$. Period j is a forecast horizon for a stability horizon s, iff*

$$1 < q_s^* \equiv q_s(i_1) = q_s(i_2) = \cdots = q_s(i_m). \tag{9}$$

If (9) holds, then $q_s^* - 1$ is a planning horizon equal to or larger than the stability horizon s.

To detect a minimal forecast horizon for a given stability horizon $s \geq 2$, the procedure described in subsection 1.2 may be adopted, merely replacing the recursive computation of the q-values via (3) by that via (8). The resulting integrated algorithm has the *same* complexity ($O(n \log n)$) as that in subsection 1.2.

1.6. Numerical Results

In this subsection we describe the results of a numerical study conducted to evaluate the magnitude of minimal forecast horizons in a variety of lot sizing problems. We have evaluated a total of 435 problems. Their one-period parameters d_i, K_i, c_i and h_i are generated from the following first-order autoregressive equations:

$$d_1 = e_1^d, \quad K_1 = e_1^K, \quad c_1 = e_1^c, \quad h_1 = e_1^h;$$

and for $i > 1$:

$$d_i = \alpha d_{i-1} + (1-\alpha)e_i^d;$$
$$K_i = \alpha K_{i-1} + (1-\alpha)e_i^K;$$
$$c_i = \alpha c_{i-1} + (1-\alpha)e_i^c;$$
$$h_i = \alpha h_{i-1} + (1-\alpha)e_i^h;$$

where the random variables in each of the sequences: $\{e_i^d: i = 1, \ldots, \infty\}$, $\{e_i^K: i = 1, \ldots, \infty\}$, $\{e_i^c: i = 1, \ldots, \infty\}$, and $\{e_i^h: i = 1, \ldots, \infty\}$ are independent and uniformly distributed on the integer values of a prespecified interval. The random variables e_i^c and e_i^h are, in all problem instances, uniformly distributed on the integer values of the interval [1, 5]. The distribution of the e_i^K and e_i^d variables differ per problem category, as specified below. The autoregressive patterns reflect correlations between consecutive parameter values, as typically observed in most practical settings. Note also that this specification corresponds with one of the most popular forecasting techniques, i.e., the technique of exponential smoothing. The extreme case where $\alpha = 0$ corresponds with fully independent and random parameter values; in the other extreme case where $\alpha = 1$, the parameter values stay completely constant over time.

Table II specifies 25 problem categories, each consisting of 15 problem instances. The problem categories differ by the assumed values of α and the intervals from which the e_i^K and e_i^d random variables are generated (uniformly). For each problem category, we report the minimum, maximum, and median values of the computed minimal forecast horizons and the associated planning horizons. The last three columns in the table specify the minimum, maximum, and median values of the total number of setups (including period 1's setup) employed over a horizon which is equal to the minimal forecast horizon.

We conclude that the minimum forecast horizon is surprisingly short in virtually all of the above problem instances. This is particularly apparent when considering the "number of setups" columns where the median values are equal to 2 in the first 20 categories and equal to 3 in the remaining 5 categories (all with $\alpha = 0.8$). In other words, a forecast horizon tends to exist even when planning over a horizon that contains no more than one or two order cycles! It is fair to surmise that in practice virtually all dynamic lot sizing models are solved over horizons that contain more than two order cycles. Note that the median values of the obtained planning horizons tend to be half as large as the corresponding forecast horizons.

The forecast and planning horizons increase almost invariably with α. In other words, the larger the interperiod variability of the parameter values, i.e., the more the dynamic lot sizing model differs from the standard EOQ model, the shorter a minimal forecast horizon can be expected! This empirical phenomenon may be understood as follows: The larger the interperiod variability of the parameter values, the more likely it is that a specific choice f for the first order

Table II
The Impact of the Correlation Between Consecutive Periods on Forecast and Planning Horizon

			Forecast Horizon			Planning Horizon			Number of Setups		
α	d	K	Min.	Max.	Median	Min.	Max.	Median	Min.	Max.	Median
0.0	[1, 10]	[1, 50]	2	5	3	1	3	2	2	3	2
		[1, 100]	3	6	4	1	4	3	2	3	2
		[1, 200]	2	8	5	1	4	3	2	3	2
		[1, 500]	3	11	7	1	7	4	2	3	2
		[1, 1000]	4	12	8	2	9	4	2	3	2
0.0	[1, 5]	[1, 50]	2	6	3	1	4	2	2	2	2
		[1, 100]	2	9	6	1	5	3	2	3	2
		[1, 200]	3	11	7	1	7	3	2	3	2
		[1, 500]	2	11	8	1	8	4	2	3	2
		[1, 1000]	6	18	8	1	10	4	2	3	2
0.2	[1, 5]	[1, 50]	2	8	4	1	4	2	2	3	2
		[1, 100]	3	9	5	1	5	3	2	3	2
		[1, 200]	2	13	8	1	8	4	2	3	2
		[1, 500]	6	17	10	3	10	7	2	3	2
		[1, 1000]	8	27	13	3	11	7	2	3	2
0.5	[1, 5]	[1, 50]	2	13	6	1	5	3	2	5	2
		[1, 100]	4	17	7	1	6	2	2	5	2
		[1, 200]	6	26	9	3	6	5	2	6	2
		[1, 500]	8	24	11	3	12	6	2	4	2
		[1, 1000]	12	39	19	3	15	10	2	4	2
0.8	[1, 5]	[1, 50]	3	17	6	1	4	2	2	8	3
		[1, 100]	5	25	9	1	6	4	2	7	3
		[1, 200]	7	19	12	2	8	5	2	4	3
		[1, 500]	9	49	20	4	10	6	2	7	3
		[1, 1000]	15	48	27	5	17	12	2	5	3

period after period 1 has characteristics which are significantly different from those pertaining to adjacent periods; for example, a low fixed or per unit order cost, or large demand value. Such a period is, therefore, more likely to be optimal for all potential horizons beyond a relatively small forecast horizon.

The above conclusions continue to hold if the interperiod variability of the parameters is increased due to an increase of the variance of the e-variables. To substantiate this assertion, we have investigated four additional problem categories (again with 15 problem instances each) in which only the variance of the e_i^K variables is systematically increased (keeping everything else unaltered). See Table III in which the problem categories are ranked in increasing order of the variance of the e_i^K-variables. The median value of the minimal forecast horizons, the planning horizons, and the number of setups, all decrease from one problem category to the next, with just two minor violations of this monotonicity pattern. Note from Table II that the median values of the minimal forecast horizons tend to *increase* rather than decrease as the upper limits of the range of the e_i^K-variables is increased, leaving everything else unaltered. Such increases have two conflicting effects on the minimal forecast horizon: The increase in the *mean* of the

Table III
The Impact of the Setup Cost Variability on Forecast and Planning Horizons

			Forecast Horizon			Planning Horizon			Number of Setups		
α	d	K	Min.	Max.	Median	Min.	Max.	Median	Min.	Max.	Median
0.5	[1, 5]	[475, 525]	22	68	34	9	16	13	2	6	3
		[450, 550]	15	72	36	6	14	10	2	7	3
		[400, 600]	14	56	34	6	14	11	2	5	3
		[250, 750]	13	57	25	6	13	10	2	6	3
		[1, 1000]	12	39	19	3	15	10	2	4	2

e_i^K-variables results in an increase of the average order cycle and hence of the forecast horizon; this effect tends to dominate that resulting from the increased variance of these variables. See, however, Chand, Sethi and Proth for counterexamples for the above monotonicity.

2. WHAT TO DO IF NO APPROPRIATE FORECAST HORIZON IS FOUND

In most practical settings dynamic lot sizing models are solved over a horizon which can be expected to include more than two or even five order cycles. The numerical results in subsection 1.6 thus suggest that a forecast horizon associated with a desirable stability interval can be detected in many practical settings. In this section we propose planning procedures for those (presumably rare) cases where no appropriate forecast horizon is detected.

We first (subsection 2.1) derive bounds for the optimality gap associated with any given first period decision, i.e., for a given choice of the first order period f after period 1 ($2 \leq f \leq n$). We also discuss the asymptotic behavior of these bounds. In subsection 2.2 we describe our proposed full horizon-based planning procedure.

2.1. Bounding the Worst-Case Optimality Gap for the Initial Decision

We bound the worst-case optimality gap that may arise over some future horizon t in excess of n (the number of periods for which adequate data are available). Thus, define:

$F_f(t)$ = the minimum cost over a horizon of length t, given that period f is chosen as the first order period after period 1, $t \geq n$;

$\epsilon(f, n)$ = $\max_{t \geq n} \max \{[F_f(t) - F(t)]/F(t)$: *all* possible parameters in periods $n+1, \ldots, t\}$.

Clearly, $F(t) = F_q(t)$ for some $q = q(i)$ with $i \in \Omega(n)$, provided that for every horizon $t \geq n$ the first order period after period 1 is guaranteed to be one of the periods $2, \ldots, n$. (This qualification can be rigorously shown to hold for sufficiently large n and mild conditions with respect to the model parameters (see, for example, Proposition 1 in Federgruen and Tzur). In practice, the qualification can be assumed to hold for relatively small values of n.) We first need the lemma that follows; see the Appendix for the proof.

For all $i = 1, \ldots, n$ let $\lambda(i)$ denote an upper bound for the *first* period whose demand cannot be satisfied optimally by an order in period i or earlier. Such upper bounds are easily computable as a function of the known parameters in periods $i, i+1, \ldots, n$, for example, as in Proposition 1 of Federgruen and Tzur. Also defined $d_{it} = \infty$ if $t > n + 1$.

Lemma 2. *Fix $t > n$.*

a. *Let $2 \leq i < j \leq n$. $F_i(t) - F_j(t) \leq U_{ij}$ where*

$$U_{ij} = \begin{cases} (K_i - K_j) + [C(i) - C(1)]d_{ij} & \text{if } C(i) \leq C(j) \quad (10a) \\ K_i + [C(i) - C(1)]d_{ij} & \text{if } C(i) > C(j). \quad (10b) \end{cases}$$

b. *Let $2 \leq j < i \leq n$ and $C_{ji}^{min} = \min_{j \leq r < i} C(r)$. $F_i(t) - F_j(t) \leq U_{ij}$ where*

$$U_{ij} = \begin{cases} (K_i - K_j) + [C(1) - C_{ji}^{\min}] & \text{if } C_{ji}^{\min} \geq C(i) \quad (11a) \\ (K_i - K_j) + [C(1) - C_{ji}^{\min}] + [C(i) - C_{ji}^{\min}]\, d_{i\lambda(i)} & \text{otherwise.} \quad (11b) \end{cases}$$

c. *Let $2 \leq i < \min\{\ell: \ell \in \Omega(n)\}$. $F_i(t) - F(t) \leq U^i$ where*

$$U^i = \max_{l \in \Omega(n)} \{F_i(\ell - 1) - F(\ell - 1)\}. \quad (12)$$

Note that the expressions to the right of (10a), (10b), (11a), and (12) are all finite; the expression to the right of (11b) is infinite only if $\lambda(i) > n + 1$.

For all $i < j$, Lemma 2 thus provides a *finite* bound for the cost differences $F_i(t) - F_j(t)$ or $F_i(t) - F(t)$, which is independent of the value of t and the unknown demand and cost trajectories beyond the horizon n. This implies that f, the first order period after period 1, can always be chosen such that the cost difference $F_f(t) - F(t)$ is uniformly bounded; for example, $f = \min\{q: q = q(i)$ for some $i \in \Omega(n)\}$.

The following theorem shows that the same is true with respect to the optimality gaps $\epsilon(f, n)$.

Theorem 4. a. *For all $2 \leq f \leq n$,*

$$\epsilon(f, n) \leq [F(n)]^{-1} \max\{U_{fq}: q = q(i) \text{ for some } i \in \Omega(n)\}. \quad (13a)$$

b. *For all $2 \leq f \leq \min\{\ell: \ell \in \Omega(n)\}$,*

$$\epsilon(f, n) \leq [F(n)^{-1}]U^f. \quad (13b)$$

Proof. This is immediate from Lemma 2, the fact that $F(t) = F_q(t)$ for some $q = q(i)$ with $i \in \Omega(n)$, and the inequality $F(t) \geq F(n)$ for all $t \geq n$.

The bounds in Theorem 4 generalize those obtained by Lee and Denardo for models without speculative motives for carrying inventories. For this special case, the bounds in part b reduce to theirs. (Note, Lee and Denardo derive bounds for $[\underline{F}_f(t) - F(t)]$ where $\underline{F}_f(t) \leq F_f(t)$ denotes the minimum cost over a horizon of length t given that *some* (not necessarily the first) order is placed in period f.) Note that for any $f \geq 2$, the upper bound in part a can be computed in $O(n)$ time. Repeating the algorithm in Federgruen and Tzur by starting in period f rather than in period 1, one obtains the values $F_f(\ell)$ $(\ell = f, \ldots, n)$ and hence U^f; its computation thus requires $O(n \log n)$ time in the most general case and $O(n)$ in all special cases identified in Federgruen and Tzur. The bound in part a is clearly simpler to compute than that of part b; we expect the latter to be tighter, in general.

Let $\epsilon_n = \min_f \in (f, n)$ denote the minimum worst-case optimality gap at stage n $(n \geq 1)$. Under weak assumptions with respect to the single period parameters $\{d_i, K_i, h_i, c_i\}$ it is, in fact, easy to show that the minimum optimality gap ϵ_n decreases to zero, as n tends to infinity. In other words, an *asymptotically* optimal first period decision can be determined.

Corollary 2. *Assume that there exists an integer $M \geq 1$, and constants h_*, c_*, K^* and c^* such that $(d_i + \cdots + d_{i+M}) \geq 1$, $h_i \geq h_*$, $c_i \geq c_*$, $K_i \leq K^*$ and $c_i \leq c^*$ for all $i = 1, \ldots, n$. Then, $\lim_{n\to\infty} \epsilon_n = 0$.*

Proof. See the Appendix.

For settings where the horizon n may be varied Corollary 2 implies that for *every* precision $\epsilon > 0$, an ϵ-optimal first order period after period 1 (f) may be found for sufficiently large values of n (i.e., there exist choices of f such that $\epsilon(f, n) \leq \epsilon$ for all n sufficiently large).

2.2. What To Do If No Appropriate Forecast Horizon is Found: A Recommended Menu of Planning Procedures

We now discuss what planning procedures should be adopted in the (presumably rare) cases where no appropriate forecast horizon is detected with the procedures described in Section 1 (see subsection 1.6). We first recommend computing the bounds for the optimality gap $\epsilon(f, n)$ (as specified in Theorem 4) for some or all values of $f \in \{q: q = q(i)$ for some $i \in \Omega(i)\}$ in the hope that this worst-case bound be acceptably small for at least one of the candidate values of f.

Recall that the set $\{q: q = q(i)$ for some $i \in \Omega(n)\}$ is a singleton if n is a forecast horizon. If n fails to be a forecast horizon, the set is likely to be very small. In any case, its cardinality is no larger than $|\Omega(n)|$. In all of our computational work we have found the latter never to be larger than 5 (see Federgruen and Tzur). It is uniformly bounded under mild parameter conditions (see Proposition 1 in Federgruen and Tzur).

This leaves us with the question of how to proceed if *every* possible first period decision is associated with a significant worst-case optimality gap or bound thereof. It is in these cases that a new optimal schedule, obtained in a rolling horizon procedure by exact optimization over the most recent horizon, may differ markedly from the previous such schedule upon which some plans may already have been based. An *equal* degree of stability, however, is achieved whether the first period decision f is fixed on the basis of one of the many history-based heuristics, or if it is fixed, for example, in accordance with the full (n-period) horizon solution, i.e., $f = q(\ell(n))$. Moreover, the latter alternative represents a superior tradeoff of current and future cost and demand considerations over *all* periods for which adequate data prevail; this is particularly true because all history-based heuristics assume constancy in all or some of the cost parameters and are therefore ill equipped to deal with the general lot sizing model.

Note, moreover, that *any* of the periods in $\{q: q = q(i)$ for some $i \in \Omega(n)\}$ is a sensible choice for the first order period decision f, *not* just the specific choice $f = q(\ell(n))$. In selecting *among* the periods in $\{q: q = q(i)$ for some $i \in \Omega(n)\}$ it is useful to be guided by the bounds for the worst-case optimality gaps, as obtained in Theorem 4.

This suggests the following full horizon-based heuristic.

Full Horizon-Based Heuristic

STEP 1. Apply the forward algorithm in Federgruen and Tzur with the above-described associated procedure for detecting minimal forecast horizons. If a minimal forecast horizon with associated planning horizon $q^* - 1$ is detected, then set $f^* = q^*$ and terminate. Otherwise, go to Step 2.

STEP 2. Use (13a), (13b), or both to compute an upper bound for $\epsilon(f, n)$ for all $f \in \{q: q = q(i)$ for some $i \in \Omega(n)\}$ and select a value of f which *minimizes this upper bound.*

APPENDIX

Proof of Lemma 2

Part a. Consider a schedule π^1 under which period j is the first order period after period 1 and with cost

$F_j(t)$. Let s(j) denote the next order period under this schedule, if any; otherwise $s(j) = t + 1$.

Case 1. $(C(i) \leq C(j))$ Let $\hat{F}(t)$ denote the cost of the schedule π^2 which orders in periods $1, i$, and $s(j)$, and follows schedule π^1 thereafter. Clearly, $F_i(t) - F_j(t) \leq \hat{F}(t) - F_j(t)$. The difference in setup costs incurred under schedules π^2 and π^1 is clearly given by $(K_i - K_j)$. For all units sold prior to period i and from period $s(j)$ on, the same variable order and holding costs are incurred under both schedules. Finally, it is easy to verify after some algebra that the differences in variable costs incurred for the units sold in the intervals $[i, j - 1]$ and $[j, s(j) - 1]$ are given by $[C(i) - C(1)]d_{ij}$ and $[C(i) - C(j)]d_{js(j)}$, respectively. Combining all cost components we conclude that $F_i(t) - F_j(t) \leq \hat{F}(t) - F_j(t) \leq (K_i - K_j) + [C(i) - C(1)]d_{ij} + [C(i) - C(j)]d_{js(j)} \leq (K_i - K_j) + [C(i) - C(1)]d_{ij}$, where the last inequality follows from the assumption that $C(i) \leq C(j)$.

Case 2. $(C(i) > C(j))$ Let $\hat{F}(t)$ denote the cost of the schedule π^2 which orders in periods 1, i, and j, and follows schedule π^1 thereafter. Clearly, $F_i(t) - F_j(t) \leq \hat{F}(t) - F_j(t)$. The difference in setup cost incurred under schedules π^2 and π^1 is clearly K_i. For all units sold prior to period i and from period j on, the same variable order and holding costs are incurred under both schedules. The difference in variable costs incurred for the units sold in the interval $[i, j - 1]$ is given by $[C(i) - C(1)]d_{ij}$, so that $F_i(t) - F_j(t) \leq \hat{F}(t) - F_j(t) \leq K_i + [C(i) - C(1)]d_{ij}$.

Part b. Let π^1 denote the schedule defined in part a. Let $a(i) = \min\{\ell \geq i: \ell \text{ is an order period under } \pi^1\}$. Let $\hat{F}(t)$ denote the cost of the schedule π^2 which orders in periods 1, i, and $a(i)$, and follows π^1 thereafter. As before, $F_i(t) - F_j(t) \leq \hat{F}(t) - F_j(t)$. The difference in setup costs incurred under schedules π^2 and π^1 is clearly bounded by $(K_i - K_j)$. (The difference equals $(K_i - K_j)$ only if under π^1 no orders are placed between periods j and i.) For all units sold prior to period j and after period $a(i)$ both schedules incur the same variable order and holding costs. Finally, it is easy to verify that the difference in variable costs incurred for the units sold in the intervals $[j, i - 1]$ and $[i, a(i) - 1]$ are bounded by $[C(1) - C_{ji}^{\min}]d_{ji}$ and $[C(i) - C_{ji}^{\min}]d_{ia(i)}$, respectively. Adding all cost components we conclude that $F_i(t) - F_j(t) \leq \hat{F}(t) - F_j(t) \leq (K_i - K_j) + [C(1) - C_{ji}^{\min}]d_{ji} + [C(i) - C_{ji}^{\min}]d_{ia(i)}$. Clearly $0 \leq d_{ia(i)} \leq d_{i\lambda(i)}$, hence (11a) and (11b).

Part c. Consider an optimal schedule π for the t-horizon problem. In view of Lemma 1, there exists a period $\ell^* \in \Omega(n)$ in which an order is placed. Let $F_i(\ell^*, t)$ denote the minimum cost of any schedule π' which places its first order in i and a subsequent order in ℓ^*. Note that from period ℓ^* on, identical costs are incurred under π and π'. Thus,

$$F_i(t) - F(t) \leq F_i(\ell^*, t) - F(t) = F_i(\ell^* - 1) - F(\ell^* - 1).$$

Proof of Corollary 2. The parameter conditions in this corollary guarantee that $\lambda(1)$ is uniformly bounded in n (see Proposition 1 in Federgruen and Tzur). Clearly, $F(n) \geq \alpha n$ for $\alpha = c_*/M$ and *all* $n \geq 1$. This implies the existence of a uniform upper bound for optimal first order periods after period 1, i.e., the existence of a number I such that for *all* $n \geq 1$, $q \leq I$ if $q = q(i)$ for some $i \in \Omega(n)$. Let

$$f(n) = \min\{q: q = q(i) \text{ for some } i \in \Omega(n)\}. \quad (14)$$

It follows from (14) that $\epsilon_n \leq \epsilon(f(n), n) \leq \alpha^{-1} n^{-1} \max\{U_{ij}: i < j \leq I\}$. The corollary follows easily because $U_{ij} < \infty$ is *independent* of n for fixed $i < j$.

ACKNOWLEDGMENT

We are indebted to Daniel Bienstock for several helpful suggestions.

REFERENCES

AGGARWAL, A. AND J. K. PARK. 1993. Improved Algorithms for Economic Lot Size Problems. *Opns. Res.* **41,** 549–571.

AXSÄTER, S. 1982. Worst Case Performance for Lot Sizing Heuristics. *Eur. J. Opnl. Res.* **9,** 339–343.

AXSÄTER, S. 1985. Performance Bounds for Lot Sizing Heuristics. *Mgmt. Sci.* **31,** 634–640.

BAKER, R. K. 1989. Lot-Sizing Procedures and a Standard Data Set—A Reconciliation of the Literature. *J. Mfg. and Opns. Mgmt.* **2,** 199–221.

BENSOUSSAN, A., J. CROUHY AND J. M. PROTH. 1983. *Mathematical Theory of Production Planning.* Advanced Series in Management, North-Holland, Amsterdam.

BENSOUSSAN, A., J. M. PROTH AND M. QUEYRANNE. 1991. A Planning Horizon Algorithm for Deterministic Inventory Management With Piecewise Linear Concave Costs. *Naval Res. Logist.* **38,** 729–842.

BITRAN, G., T. L. MAGNANTI AND H. H. YANASSE. 1984. Approximation Methods for the Uncapacitated Dynamic Lot Size Problem. *Mgmt. Sci.* **30,** 1121–1140.

CHAND, S. 1982. Lot Sizing for Products With Finite Demand Horizon and Periodic Review Policy. *Eur. J. Opnl. Res.* **11,** 145–148.

CHAND, S., AND T. E. MORTON. 1986. Minimal Forecast Horizon Procedures for Dynamic Lot Size Models. *Naval Res. Logist. Quart.* **33,** 111–122.

CHAND, S., S. SETHI AND J. M. PROTH. 1990. Existence of Forecast Horizons in Undiscounted Discrete-Time Lot Size Model. *Opns. Res.* **38,** 884–892.

CHAND, S., S. SETHI AND G. SORGER. 1992. Forecast Horizons in the Discounted Dynamic Lot Size Model. *Mgmt. Sci.,* **38,** 1034–1048.

DIXON, P. S., AND E. A. SILVER. 1981. A Heuristic Solution Procedure for the Multi-Item, Single Level, Limited Capacity Lot Sizing Problem. *J. Opns. Mgmt.* **2,** 23–39.

EPPEN, G. D., F. J. GOULD AND B. P. PASHIGIAN. 1969. Extensions of Planning Horizon Theorem in the Dynamic Lot Size Model. *Mgmt. Sci.* **15,** 268–277.

FEDERGRUEN, A., AND M. TZUR. 1991. A Simple Forward Algorithm to Solve General Dynamic Lot Sizing Models With n Periods in $O(n \log n)$ or $O(n)$ Time. *Mgmt. Sci.* **37,** 909–925.

FLORIAN, M., J. K. LENSTRA AND A. H. G. RINNOOY KAN. 1980. Deterministic Production Planning: Algorithms and Complexity. *Mgmt. Sci.* **26,** 12–20.

KARNI, R. 1981. Maximum Pert Period Gain (MPG)—A Lot Sizing Procedure for Unconstrained and Constrained Requirements Planning Systems. *Prod. and Inv. Mgmt. 22,* 91–98.

LEE, C., AND E. DENARDO. 1986. Rolling Planning Horizons: Error Bounds for the Dynamic Lot Size Model. *Math. Opns. Res.* **11,** 423–432.

LUNDIN, R. A., AND T. E. MORTON. 1975. Planning Horizons for the Dynamic Lot Size Model: Zabel vs. Protective Procedures and Computational Results. *Opns. Res.* **23,** 711–734.

NAHMIAS, S. 1989. *Production and Operations Analysis.* Irwin, Horewood, Ill.

PETERSON, R., AND E. SILVER. 1979. *Decision Systems for Inventory Management and Production Planning*. John Wiley, New York.

SILVER, E. A., AND H. C. MEAL. 1973. A Heuristic for Selecting Lot-Size Quantities for the Case of a Deterministic Time-Varying Demand Rate and Discrete Opportunities for Replenishment. *Prod. and Inv. Mgmt.* **14,** 64–74.

STEELE, D. C. 1973. The Nervous MRP Systems: How to do Battle. *Prod. and Inv. Mgmt.* **16,** 83–89.

VAN HOESEL, S., AND A. WAGELMANS. 1993. Sensitivity Analysis of the Economic Lot-Sizing Problem. *Discr. Appl. Math.* **45,** 291–312.

WAGELMANS, A., S. VAN HOESEL AND A. KOLEN. 1992. Economic Lot-Sizing: An $O(n \log n)$-Algorithm That Runs in Linear Time in the Wagner-Whitin case. *Opns. Res.* **40,** S145–S156.

WAGNER, H. M., AND T. M. WHITIN. 1958. Dynamic Version of the Economic Lot Size Model. *Mgmt. Sci.* **5,** 89–96.

ZABEL, E. 1964. Some Generalization of Inventory Planning Horizon Theorem. *Mgmt. Sci.* **10,** 465–471.