THE JOINT REPLENISHMENT PROBLEM WITH TIME-VARYING COSTS AND DEMANDS: EFFICIENT, ASYMPTOTIC AND ε-OPTIMAL SOLUTIONS

AWI FEDERGRUEN

Columbia University, New York, New York

MICHAL TZUR

Tel Aviv University, Tel Aviv, Israel

We address the Joint Replenishment Problem (JRP) where, in the presence of joint setup costs, dynamic lot sizing schedules need to be determined for m items over a planning horizon of N periods, with general time-varying cost and demand parameters. We develop a new, so-called, partitioning heuristic for this problem, which partitions the complete horizon of N periods into several relatively small intervals, specifies an associated joint replenishment problem for each of these, and solves them via a new, efficient branch-and-bound method. The efficiency of the branch-and-bound method is due to the use of a new, tight lower bound to evaluate the nodes of the tree, a new branching rule, and a new upper bound for the cost of the entire problem. The partitioning heuristic can be implemented with complexity $O(mN^2\log\log N)$. It can be designed to guarantee an ε-optimal solution for any $\epsilon > 0$, provided that some of the model parameters are uniformly bounded from above or below. In particular, the heuristic is asymptotically optimal as $N \to \infty$ for any fixed number of items m, and it remains asymptotically optimal when both m and N are simultaneously increased to infinity. Most importantly, a numerical study shows that the partitioning heuristic performs exceptionally well. Even for small problems, the average optimality gap is only 0.38% and in none of the problem categories is it larger than 0.78%.

This paper addresses the Joint Replenishment Problem (**JRP**), one of the most basic multi-item lot sizing models. In the **JRP**, the need to coordinate planning for the different items arises from the potential to exploit economies of scale when placing orders for more than one item at a time.

The model specifies a horizon divided into a finite number of (say N) periods, each with a known demand for m items, all of which must be satisfied. An unlimited amount may be ordered (produced) in each period. The cost structure consists of fixed-order costs, variable order costs which are proportional with the order sizes, and holding costs which are proportional with the end-of-the-period inventory levels. The fixed-order costs consist of a joint setup cost incurred for *any* order, regardless of its composition, and item-specific setup costs for any specific item included in the order. Demands and all cost parameters may vary over time. The objective is to determine an ordering policy that will minimize total costs while satisfying all demands without backlogging.

Arkin, Joneja and Roundy (1989) proved that the **JRP** is NP-complete. Indeed, existing exact solution methods are of a complexity which grows exponentially with the number of periods or the number of items, and therefore may be used only for problems of small size. Even for the case where *all* parameters are constant over time, no polynomial-time solution methods are known. Indeed, at least since Brown (1967), attention has focused on heuristics, primarily for the case where all cost and demand parameters are constant over time, and, more recently, for settings that allow for time-varying demands (while continuing to assume constant cost parameters).

We first develop an exact branch-and-bound method for the **JRP** with time-varying parameters. This method is efficient for problems of moderate size, i.e., it is capable of solving problems with 20–30 periods and 20–30 items in a reasonable amount of time (see Section 2).

The efficiency of the branch-and-bound method is due to the use of a new, tight lower bound to evaluate

the nodes of the tree, a new branching rule, as well as a new upper bound for the cost of the entire problem. To obtain the upper bound, we formulate the **JRP** as a problem of minimizing a supermodular set function, apply a greedy procedure to this set minimization problem, and use the associated cost value. This greedy heuristic comes close to being optimal as we demonstrate both numerically and on the basis of a worst-case analysis.

The results on forecast horizons for single item, dynamic lot sizing models, reported in Federgruen and Tzur (1994), suggest that optimal or close-to-optimal initial decisions can be made by truncating the horizon after a relatively small number of periods. In all instances of the problem investigated there, a forecast horizon is found in which at most three and, usually only two, orders are placed (the obligatory order in the first period included). It is reasonable to expect that similarly short forecast horizons continue to prevail in the presence of joint setup costs. This, together with the attractiveness of our branch-and-bound procedure for relatively short problem horizons, suggests that a close-to-optimal solution may be generated by partitioning the horizon into smaller intervals.

Therefore, we develop a new heuristic for the **JRP**, which we refer to as the partitioning heuristic. This heuristic partitions the complete horizon of N periods into several intervals, and specifies an associated **JRP** for each of these. The lengths of the intervals permit the use of the exact branch-and-bound method mentioned above. The **JRP** associated with a given interval is obtained from the restriction of the complete problem to the periods of that interval, with an option to choose starting conditions that appropriately complement the solutions obtained in prior intervals.

The partitioning heuristic can be implemented with complexity $O(mN^2 \log\log N)$, where m denotes the number of items and N the length of the horizon. It can be designed to guarantee an ϵ-optimal solution for any $\epsilon > 0$, provided that some of the model parameters are uniformly bounded from above or *from* below. In particular, the heuristic is asymptotically optimal as $N \rightarrow \infty$ for any fixed number of items m, and it remains asymptotically optimal when both m and N are simultaneously increased to infinity. Most importantly, a numerical study reported in Section 3 shows that the partitioning heuristic performs exceptionally well. Even with very small intervals, the average optimality gap, when measurable, is only 0.38% and in none of the problem categories is it larger than 0.78%.

Exact solution methods which are based on dynamic programming (with exponential complexity) have been proposed by Zangwill (1966), Veinott (1969), Kao (1979), and Silver (1979). These exploit the zero inventory ordering property, which states that there exists an optimal solution in which no order is placed for any item unless its inventory is down to zero (see Veinott). Erenguc (1988) developed a branch-and-bound procedure which differs from ours in the choice of lower and upper bounds, as well as in its branching rules. This author reports numerical results for problems with up to 12 periods and 20 items. An exact cutting plane method was recently developed by Raghavan and Rao (1991); these authors report numerical results for problems with up to 30 periods and 20 items. A more recent version (Raghavan and Rao 1992) reports on problems with up to 40 periods and 100 items (see Section 2 for a detailed discussion).

As mentioned, most of the proposed heuristic methods deal with the special case where all cost and demand parameters are constant over time. Jackson, Maxwell and Muckstadt (1985) showed that in this special case, a simple, so-called, power-of-two policy can be constructed whose long-run average cost comes within 6% of the minimum cost value. Roundy (1985) showed for a more general model, i.e., the one warehouse multiretailer problem, that the worst-case optimality gap for power-of-two policies is, in fact, limited to 2% only. Federgruen and Zheng (1992) extended the 2% optimality gap result to joint replenishment problems with general (submodular) joint cost structures. We refer the reader to the latter paper and to Aksoy and Erenguc (1988) for a detailed review of previous heuristics for models with constant parameters. Kao (1979) and Leopoulos and Proth (1985) propose heuristics for models with time-varying demands. Both papers report results for the two-product problem; it is unclear how their complexity grows with the number of items. The optimality gap of Kao's heuristic can be arbitrarily large, as demonstrated in Joneja (1990). Silver (1976) reports on a single-pass heuristic with a three-product example.

Atkins and Iyogun (1988) extend the well-known Silver-Meal heuristic of the single item model to joint replenishment problems with time-varying demands but constant cost parameters. They also derive a lower bound for the minimum cost value which provides the foundation for ours. The only heuristics with known worst-case optimality gaps are due to Iyogun (1987, chap. 4) and Joneja (1990) which apply to models with constant cost parameters. To our knowledge, our heuristic is the first with a bounded

worst-case optimality gap for models in which both the demand and the cost parameters vary over time. (As mentioned, it is possible to design this heuristic to achieve an ϵ-optimality gap for any $\epsilon > 0$.)

The remainder of this paper is organized as follows: In Section 1 we develop our lower and upper bounds. The branch-and-bound procedure is described in Section 2. Section 3 is devoted to the partitioning heuristic.

1. LOWER AND UPPER BOUNDS

In this section, we describe lower and upper bounds for the minimum cost value z^*. These bounds may be used in the branch-and-bound procedure described in Section 2. Alternatively, they are useful in evaluating heuristics, e.g., those described in Section 3.

The **JRP** is specified by the parameters:

N = the number of periods;
m = the number of items;

For $i = 1, \ldots, m$ and $t = 1, \ldots, N$:

d_{it} = the demand for item i in period t (we assume, without loss of generality, that $d_{it} \geq 0$);
K_{it} = the fixed-order cost for item i in period t;
K_{0t} = the fixed joint order cost in period t;
c_{it} = the variable per unit order cost of item i in period t;
h_{it} = the cost of carrying a unit of inventory of item i at the end of period t;

We also assume that orders are filled instantaneously. (We can easily adjust our planning methods to the case where each order for item i arrives after ℓ_i periods, $i = 1, \ldots, m$.)

The following is a mixed integer programming formulation of the problem: For $i = 1, \ldots, m$ and $t = 1, \ldots, N$, let

X_{it} = the production volume of item i in period t;
I_{it} = the ending inventory of item i in period t;

$$Y_{it} = \begin{cases} 1 & \text{if } X_{it} > 0 \\ 0 & \text{otherwise} \end{cases}$$

$$Y_t = \begin{cases} 1 & \text{if } \sum_{i=1}^{m} X_{it} > 0 \\ 0 & \text{otherwise} \end{cases}$$

$$z^* = \min \left\{ \sum_{t=1}^{N} [K_{0t} Y_t + \sum_{i=1}^{m} K_{it} Y_{it} + \sum_{i=1}^{m} c_{it} X_{it} + \sum_{i=1}^{m} h_{it} I_{it}] \right\}$$

subject to

$$I_{it} = I_{it-1} + X_{it} - d_{it} \quad i = 1, \ldots, m, \ t = 1, \ldots, N$$

$$X_{it} \leq \left(\sum_{r=t}^{N} d_{ir} \right) Y_{it} \quad i = 1, \ldots, m, \ t = 1, \ldots, N$$

$$Y_t \geq \left(\sum_{i=1}^{m} Y_{it} \right) \bigg/ m \quad t = 1, \ldots, N$$

$$I_{i0} = I_{iN} = 0; \ X_{it} \geq 0; \ I_{it} \geq 0;$$

$$Y_{it} = 0, 1; \quad Y_t = 0, 1.$$

1.1. Lower Bounds

Atkins and Iyogun observed that a lower bound can be obtained by replacing the joint setup cost structure by one with *separable* costs, i.e., modified item specific setup costs:

$$K'_{it} = K_{it} + \alpha_{it} \quad i = 1, \ldots, m, t = 1, \ldots, N, \quad (1)$$

where

$$\sum_{i=1}^{m} \alpha_{it} = K_{0t} \quad t = 1, \ldots, N$$

$$\alpha_{it} \geq 0 \quad i = 1, \ldots, m, t = 1, \ldots, N. \tag{2}$$

(Note that for any given strategy and associated order periods, the setup cost incurred in any order period in the transformed model is no larger than the setup cost incurred in the original model.)

Let $z_i(\alpha_{i1}, \ldots, \alpha_{iN})$ denote the minimum cost for *item i* in the single item dynamic lot sizing model with modified setup costs K'_{it} specified by (1). Observe that

$$\sum_{i=1}^{m} z_i(\alpha_{i1}, \ldots, \alpha_{iN}) \leq z^* \tag{3}$$

for any collection of allocated setup costs $\{\alpha_{it}: i = 1, \ldots, m; t = 1, \ldots, N\}$ which satisfy (2). The best lower bound of this type is clearly obtained by determining:

$$\max_{\alpha: \ \alpha \text{ satisfies (2)}} \sum_{i=1}^{m} z_i(\alpha_{i1}, \ldots, \alpha_{iN}). \tag{4}$$

We also observe the following:

Lemma 1. *Each of the functions z_i, as well as the function z, is jointly concave.*

Proof. Each of the functions z_i can be represented as the minimum of a number of affine functions in the α-variables (one for each of the $2^N - 1$ possible sets of order periods).

Evaluation of (4) requires computing the maximum of a concave, but nonseparable, objective on the polyhedron described by (2). Standard concave

programming techniques may be invoked. Alternatively, (4) may be evaluated by solving the linear programming relaxation of the integer programming formulation in Joneja (1990), straightforwardly extended to allow for nonstationary cost parameters. Joneja (1989, Theorem 2.3) establishes that the optimum value of this linear program equals that of a specific Lagrangian dual and the latter reduces to (4). The linear program uses $O(mN^2)$ variables and mN constraints. No efficient exact solution method for (4) is known.

One may consider evaluating (4) by the greedy procedure. The discrete version of this procedure starts with $\alpha = 0$ and in each iteration increments by one unit an α_{it} variable for which this is feasible (i.e., does not violate (2)) and for which the resulting objective function value improvement is maximal; the procedure terminates when no α_{it} variable can be incremented. The polyhedron in (2) is a so-called polymatroid (see Edmonds 1970 or Nemhauser and Wolsey 1988); the greedy procedure results in an optimal solution for any separable concave objective, as well as certain types of nonseparable objectives. Since the objective in (4) fails to be one of these types, the greedy procedure is likely to generate a good, but not necessarily, optimal solution. (See Federgruen and Groenevelt 1986 for more details.)

The greedy heuristic is also quite expensive! It consists of $\max_t K_{0t}$ iterations; in each iteration one needs to evaluate the impact of incrementing up to mN variables, each of which requires the solution of a single-item dynamic lot sizing model. Thus, the total complexity of the greedy heuristic is $O((\max_t K_{0t})mN^2 \log N)$. We therefore describe a simpler heuristic procedure.

For each period we determine the best allocation vector $\alpha^{*t} = (\alpha_1^{*t}, \ldots, \alpha_m^{*t})$ for the joint setup cost value K_{0t}, i.e., the allocation vector which maximizes the total cost over all single-item dynamic lot sizing models that arise when all setup costs in the ith model are increased by α_i^{*t} ($i = 1, \ldots, m$). The advantage of this heuristic is that each period's allocation vector is determined separately, even though its choice continues to be based on maximization of a total cost criterion over the complete planning horizon. A second advantage is that the maximization consists of optimizing a separable concave objective, subject to a single budget constraint. Indeed, the allocation vector $\alpha^{*t} = (\alpha_1^{*t}, \ldots, \alpha_m^{*t})$ is chosen to achieve the maximum in:

$$\max_{\alpha = (\alpha_1, \ldots, \alpha_m)} \sum_{i=1}^{m} Z_i(\alpha_i)$$

subject to   (5)

$$\sum_{i=1}^{m} \alpha_i = K_{0t}; \quad \alpha_i \geq 0,$$

where for all $i = 1, \ldots, m$; $Z_i(\alpha) = z_i(\alpha, \alpha, \ldots, \alpha)$ = the minimum cost in item i's dynamic lot sizing model with setup costs $K'_{it} = K_{it} + \alpha$ for all $t = 1, \ldots, N$. (See the definition of $z_i(\cdot, \ldots, \cdot)$ above.) (Note that the optimum value of (5) may, in itself, fail to be a lower bound for z^* unless $K_{0t} = \min\{K_{0\tau}: \tau = 1, \ldots, N\}$. However, $z(\alpha^{*1}, \ldots, \alpha^{*N})$ is a lower bound by (3).)

Observe that each of the functions $Z_i(\cdot)$ is concave as well because it too can be represented as the finite minimum of a number of affine functions. Problem 5 thus amounts to determining the maximum of a separable concave objective subject to nonnegativity constraints for the individual variables and a single budget constraint. It is well known (see Gross 1956 and Fox 1966) that the greedy procedure described above terminates with an optimal solution. In other words, the vector α^{*t} determines the optimal stationary allocation of the joint setup cost value K_{0t}. We refer to the resulting lower bound as z_{LB}.

This procedure needs to be applied to the single instance of (5) with K_{0t} replaced by $K_0^{\max} \equiv \max\{K_{0\tau}: \tau = 1, \ldots, N\}$ as all other required allocation vectors α (for smaller values of K_{0t}) are obtained from intermediate iterations of this procedure. The greedy procedure requires $2m$ solutions of single-item dynamic lot sizing models to evaluate $Z_i(0)$ and $Z_i(1)$ for all $i = 1, \ldots, m$, and one additional such solution in each of the $(K_0^{\max} - 1)$ subsequent iterations. Each single-item dynamic lot sizing model is solved in $O(N \log N)$ time using one of the methods in Federgruen and Tzur (1991), Wagelmans, Van Hoesel and Kolen (1992) or Aggarwal and Park (1993). The overall complexity is thus $O((m + K_0^{\max})N \log N)$. Polynomial-time discrete algorithms for (5) are due to Galil and Megiddo (1979), Frederickson and Johnson (1982), and Groenevelt (1986) (see also Zipkin 1980).

1.2. Upper Bound (The Greedy-Add Heuristic)

We now describe the implementation of the so-called greedy-add heuristic (see Cornuejols et al. 1977) to the JRP. The value of this heuristic's solution will be used in the branch-and-bound procedure as an upper bound to the optimal cost. The greedy-add heuristic

determines the set of order periods as follows: Starting with an empty set, this set is incremented in each iteration by a single period, say period j, whose addition as an order period results in the largest cost savings, i.e. $\Phi(S \cup \{j\}) = \min_{i \notin S} \Phi(S \cup \{i\})$. The heuristic terminates when $S = \{1, \ldots, N\}$ or when the cost cannot be reduced by the addition of a period to S, i.e., $\Phi(S \cup \{j\}) \geq \Phi(S)$ for all $j \notin S$.

Note that the greedy-add heuristic is likely to give good results by itself. This follows from the following characterization of the optimality gap of the greedy-add heuristic. Let z^G denote the cost value of this heuristic.

Proposition 1

$$\frac{(z^G - z^*)}{(\Phi(\{\phi\}) - z^*)} \leq \frac{1}{e}.$$

Proof. It suffices to show that the set function $\Phi(\cdot)$ is supermodular, i.e., $\Phi(S \cup \{j\}) - \Phi(S) \leq \Phi(T \cup \{j\}) - \Phi(T)$ for all $S \subseteq T$ and $j \notin T$. The proposition then follows from Nemhauser, Wolsey and Fisher (1978) who proved it for general set minimization problems with supermodular set functions. Note that

$$\Phi(S) = \sum_{t \in S} K_{0t} + z(\alpha^1, \ldots, \alpha^N), \tag{6}$$

where the z function was defined in Section 1, $\alpha^t = (\alpha_{1t}, \ldots, \alpha_{mt})$ and for all $i = 1, \ldots, m$

$$\alpha_{it} = \begin{cases} 0 & \text{if } t \in S \\ \infty & \text{otherwise.} \end{cases}$$

Note that the first term in (6) is modular in S. The supermodularity of the last term follows from the joint concavity of the z-function (see Lemma 1).

The greedy heuristic has complexity $O(mN^3 \log N)$ because it terminates after at most N iterations and at each iteration at most N values of the function $\Phi(\cdot)$ need to be evaluated. As we show in our computational study in the next section, the greedy heuristic may be used to obtain an excellent upper bound for the problem.

2. AN EXACT BRANCH-AND-BOUND METHOD

The **JRP** clearly can be solved in $O(m2^N N \log N)$ time by enumerating all possible sets of periods in which some order is placed, and solving for each such set the resulting m independent, single-item, dynamic lot sizing problems. This is prohibitively expensive for all but small values of N.

In this section, we describe an exact branch-and-bound procedures, an *implicit* enumeration method

for the **JRP**. As our numerical results below indicate, this method can be used comfortably for problems with 20–30 periods and 20–30 items. More importantly, it is used to solve the subproblems which arise in the partitioning heuristic described in Section 3.

Each node of the branch-and-bound tree is characterized by a partition of $\{1, \ldots, N\}$ into three sets S^+, S^-, and S^0. The set S^+ includes all periods in which one is committed to incur the joint setup cost, the set S^- includes the periods in which one is committed *not* to incur the joint setup cost, and the set $S^0 = \{1, \ldots, N\} \backslash (S^+ \cup S^-)$ includes the periods in which one is not committed yet. For each node in the tree, we evaluate the lower bound

$$\sum_{t \in S^+} K_{0t} + \sum_{i=1}^m z_i(\alpha_{i1}, \ldots, \alpha_{iN}) \tag{7}$$

with

$$\alpha_{it} = \begin{cases} 0 & \text{if } t \in S^+ \\ \infty & \text{if } t \in S^- \\ \alpha_i^{*t} & \text{if } t \in S^0 \end{cases}$$

where α_i^{*t} is computed from the lower bound described in the previous section. The expression in (7) provides a lower bound in view of (3). It requires m solutions of single-item lot sizing models. Each of these lower bounds is associated with a feasible collection of schedules for the m items. The cost of the feasible solution provides an upper bound for z^* and is easy to obtain by adding to (7) the *unaccounted* part of the joint setup costs in the periods in S^0 in which some order is placed. At any stage of the branch-and-bound procedure, the best available upper bound may be used to eliminate parts of the tree.

Every node in the tree has two successor nodes; the first (second) successor node has an additional period in S^0 shifted to S^- (S^+). Therefore, the branch-and-bound procedure is completely specified by the choice of the branching rule, i.e., a rule to select a period from the set S^0. Our proposed branching rule employs a fixed ranking of the periods $\{1, \ldots, N\}$ and chooses the highest ranked period in S^0.

2.1. The Branching Rule

Rank the periods $t \in \{1, \ldots, N\}$ in increasing order of the value

$$\sum_{i=1}^m d_{it} h_{it-1} - \sum_{i=1}^m c_{it} \sqrt{\frac{2 \bar{K}_i \bar{d}_i}{\bar{h}_i}} - \left[\sum_{i=1}^m K_{it} + K_{0t} \right], \tag{8}$$

where $\bar{K}_i$, $\bar{h}_i$, $\bar{d}_i$ are averages over the N periods of item i's individual setup cost, holding cost, and demand, respectively. This value represents a proxy for

the potential savings obtained when deciding on a joint setup in period t, and is based on the assumption that all items are included in the joint order. The first term represents the savings in holding costs at the end of period $t - 1$, i.e., at the beginning of period t. The second term denotes the variable order costs in the same period t, assuming that the items' order sizes in this period are chosen as if all cost and demand parameters are stationary and equal to their average value over the planning horizon, i.e., assuming the order sizes equal the optimal order quantities in the corresponding EOQ models. The third and final term denotes the total setup cost of this period (the remaining cost component of period t). Thus, the index value of (8) is a measure for the *relative* attractiveness of a joint setup in a given period. Note that the branching rule only depends on the *relative* ranking of the period measure in (8).

In the proposed branching rule, any element in S^0 with the lowest index value is shifted to S^- and S^+ when specifying the two successor nodes for a given node in the branch-and-bound tree. In our proposed depth-first method, we investigate the first successor node and its subtree before turning to the second successor node.

We have investigated several other branching rules which almost always led to inferior results to those reported here. (See Tzur 1992, Section 7.2 for more details.)

2.2. Computational Study of the Branch-and-Bound Procedure

We evaluated the performance of this branch-and-bound method with respect to a collection of hundreds of problem instances, partitioned into several sets. These problem sets are generated from a basic class of problem instances by systematically varying one, or sometimes two of the parameter sets. This basic class of instances has a planning horizon of $N = 18$ periods, and $m = 5$ items. Its cost and demand parameters are generated from the first-order autoregressive equations:

$$d_{i1} = e_{i1}^d, \quad c_{i1} = e_{i1}^c, \quad h_{i1} = e_{i1}^h \quad (i = 1, \ldots, m);$$

$$K_{i1} = e_{i1}^K \quad (i = 0, \ldots, m);$$

and for $t > 1$:

$$d_{it} = \alpha d_{i,t-1} + (1 - \alpha)e_{it}^d;$$

$$c_{it} = \alpha c_{i,t-1} + (1 - \alpha)e_{it}^c;$$

$$h_{it} = \alpha h_{i,t-1} + (1 - \alpha)e_{it}^h;$$

$$K_{it} = \alpha K_{i,t-1} + (1 - \alpha)e_{it}^K;$$

where the sequences: $\{e_{it}^d: i = 1, \ldots, m, t = 1, \ldots, N\}$, $\{e_{it}^c: i = 1, \ldots, m, t = 1, \ldots, N\}$, $\{e_{it}^h: i = 1, \ldots, m, t = 1, \ldots, N\}$, and $\{e_{it}^K: i = 0, \ldots, m, t = 1, \ldots, N\}$ are independent random variables, uniformly distributed on the integer values of a prespecified interval. The autoregressive patterns reflect correlations between consecutive parameter values, as typically observed in most practical settings. 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.

In practice, one tends to find that cost parameters are more stable than demands. On the other hand, the opposite relationship may apply when, for example, a significant part of the variable cost rates consist of widely fluctuating costs of raw materials, energy sources, or seasonally available labor. We investigate the performance for five distinct values of α, the adjustment factor (the extreme values $\alpha = 0$ and $\alpha = 1$ included). The four autoregressive equations which describe the dynamics of the demand and cost parameters have been specified to employ the same value for α, even though, in most practical settings, equation-specific values of α would apply. Note, however, that scenarios generated with a high value for α and large ranges for the parameter values, represent settings with some items consistently more expensive (in one or more cost parameters) than others. Scenarios with a high value for α and small ranges for the parameter values, represent settings in which the items are virtually identical and with small parameter variations over time. Finally, scenarios with a low value for α and large ranges for the parameter values represent settings with large nonstationarities but with none or few of the items systematically more expensive than others. (Systematic variation of equation-specific α-values would greatly increase the number of problem instances.)

In our basic class of instances, $\alpha = 0.5$ and the random variables e_{it}^d, e_{it}^c, e_{it}^h and e_{it}^K ($i = 1, \ldots, m$, $t = 1, \ldots, N$) are uniformly distributed on the integer values of the intervals [1, 10], [5, 10], [1, 5], and [10, 30], respectively. The random variables e_{0t}^K ($t = 1, \ldots, N$) are uniformly distributed over the integer values of the interval [80, 120].

In problem set 1 we vary the horizon length N from $N = 12$ to $N = 36$, and in problem set 2 we vary the number of items from $m = 5$ to $m = 35$. In the third problem set we vary the number of items from $m = 10$ to $m = 20$ for a horizon length $N = 24$. In problem set 4 we vary the value of α, i.e., the degree of correlation between consecutive parameter values

between 0 and 1. Finally, in the rest of the problem sets we vary the ranges of the demand and cost parameters (one at a time). Each problem set consists of 10 independently generated instances for each of the considered combinations of N, m, α, and demand and cost parameters.

We executed the branch-and-bound method on an IBM 4381 in FORTRAN. Tables I–IV report on the method's performance with regard to the four problem sets, respectively. We report both the CPU times in seconds, and the number of nodes investigated in the branch-and-bound tree. The numbers reported within parentheses express the number of nodes as a percentage of the 2^{N-1} possible nodes. (Since the first period demand is positive in all of the considered problem instances, there are 2^{N-1} distinct sets of periods in which a joint setup cost is incurred.) We characterize in addition the performance of the greedy heuristic and the lower bound z_{LB}, reporting the values z^G/z^* and z^*/z_{LB}.

We are pleased to conclude that the branch-and-bound method is capable of solving in a reasonable amount of time problems with as many as 20–30 periods and 20–30 items.

Tables I–IV exhibit that the greedy heuristic by itself is extremely close to optimal, with an overall average optimality gap of 0.47%; moreover, the average optimality gap is always less than 1.2% for all considered combinations of N, m, and α. The lower bound z_{LB} is relatively accurate as well. For the base value of α ($\alpha = 0.5$) the average accuracy

gap is remarkably stable across all considered combinations of N and m, varying between 1.9% and 4.5% with the vast majority in the 2.4–2.8% range. The accuracy gap is, however, much more sensitive to the value of α, as exhibited in Table IV. Large (small) values of α, i.e., a higher (lower) degree of intertemporal correlation, result in a significantly lower (higher) accuracy gap: e.g., for $\alpha = 1$ (constant parameters) the average accuracy gap is only 0.07%: This case is comparable to the stationary model with the long-run average cost criterion in which even the worst-case accuracy gap of this lower bound is 6%, but the empirically observed average is one or two magnitudes smaller (see Jackson, Maxwell and Muckstadt). For $\alpha = 0$, i.e., when the parameters are completely independent over time, the average accuracy gap is 7%. The increased gap is likely to be due to the stationary allocation heuristic discussed in Section 1, exhibiting a more considerable optimality gap vis-à-vis the best possible lower bound in (4).

The branch-and-bound tree requires the investigation of a small percentage of the possible number of nodes, only. Moreover, this percentage decreases rapidly with the horizon length N, i.e., by a factor of about 10 each time the horizon length is increased by six. Problems with up to 24 periods and 5 items can be solved to optimality in less than a CPU minute; problems with 36 periods and 5 items require about 20 CPU minutes on average.

Table I

Problem Set 1, $m = 5$

N	12	18	24	30	36
Average CPU time	0.4	4.0	40.1	231.7	1,401.2
Average number of nodes	90	378.8	2,855.2	12,595.2	74,774.4
	(4.4%)	(0.29%)	(0.03%)	(0.002%)	(0.0002%)
z^G/z^*	1.0124	1.0038	1.012	1.0054	1.0071
z^*/z_{LB}	1.0447	1.0254	1.0268	1.0254	1.0278

Table II

Problem Set 2, $N = 18$

m	5	10	15	20	25	30	35
Average CPU time	4.0	19.4	48.0	116.3	195.6	149.3	87.9
Average number of nodes	378.8	973.2	1,678.6	3,073.4	4,100.4	2,744.2	1,377.6
	(0.29%)	(0.7%)	(1.3%)	(2.3%)	(3.1%)	(2.1%)	(1.1%)
z^G/z^*	1.0038	1.0047	1.0025	1.0018	1.00057	1.00012	1.0
z^*/z_{LB}	1.0254	1.026	1.0244	1.0252	1.0245	1.0217	1.0186

Table III

Problem Set 3, $N = 24$

m	10	15	20
Average CPU time	250.4	727.1	1,402.5
Average number of nodes	9,163.2	19,396.8	34,108.0
	(0.11%)	(0.23%)	(0.41%)
z^G/z^*	1.0097	1.0065	1.0034
z^*/z_{LB}	1.0279	1.0276	1.0265

Tables II and III exhibit that the computational effort is quite sensitive to the number of items involved, even when measured by the average number of nodes evaluated in the branch-and-bound tree. (Recall that the computational effort involved in evaluating a *single* node is roughly proportional to the number of items.) Observe, for example, in Table II that the percentage of nodes evaluated increases when the number of items is increased from $m = 5$ to $m = 25$, even though the quality of both the upper and the lower bounds is *insensitive* to the number of items. This increase is probably due to our chosen branching rule which applies depth-first and shifts periods, characterized by the index (8) as unlikely setup periods, into S^- first; problems with a *low* number of items have a relatively small number of setup periods so that the periods *initially* shifted into S^- are more likely to be classified correctly as periods without setups, compared to problems with a large number of items. Indeed, it might be desirable to choose the branching rule as a function of the "expected" cardinality of the set S^+ in the optimal solution.

Observe also that the percentage of nodes evaluated *decreases* when the number of items is larger than 25. Here we observe that the accuracy gap of the upper bound, small as it already is when $m = 25$, decreases by more than a factor of four when m is increased from 25 to 30, and to zero when m is increased from 30 to 35. This decrease is due to both the optimal solution and the greedy heuristic prescribing setups in all but an ever decreasing, small number of periods. (For $m = 35$, setups are prescribed in all periods across all 10 problem instances: Hence, the 0% optimality gap.)

We conclude that for fixed N, and under our chosen branching rule, the number of investigated nodes in the branch-and-bound tree varies as a function of the relative sparsity of the optimal collection of setup periods S^+ in $\{1, \ldots, N\}$. If the optimal S^+ is very sparse or dense in $\{1, \ldots, N\}$ the percentage of investigated nodes is low, and it is highest for intermediate levels of sparsity. The same phenomenon would thus arise if K_0 were increased along with m, with the number of evaluated nodes varying with m as a unimodal function, reaching a similar maximum value for a larger maximizing value of m.

Given the prevalence of a large number of extremely close-to-optimal nodes, we conjecture that a significant reduction in computational effort can be achieved if the branch and bound procedure is terminated as soon as a solution is found to be within a given reasonable gap ϵ of the optimal value.

Table IV indicates that the computational requirements of the branch-and-bound tree vary significantly with the value of α. This phenomenon can be explained by the above-discussed variations in the accuracy of the lower bound.

In the rest of the problem sets we investigated the performance of the branch-and-bound procedure with respect to changes in the ranges of the demand and cost parameters. We conclude that the results are insensitive to those changes (see Tzur, Section 7.2).

The best alternative exact solution method appears to be Raghavan and Rao's (1991) cutting plane algorithm. They implemented their algorithm on a Sun 490 Server, and tested problems ranging in size from 3 items and 12 periods to problems with 20 items and 30 periods. These problems were solved to optimality with CPU times ranging between 2.75 seconds for the smallest problem to 5,054 seconds for the largest

Table IV

Problem Set 4, $N = 18$, $m = 5$

α	0.0	0.2	0.5	0.8	1.0
Average CPU time	25.1	13.2	4.0	1.2	1.1
Average number of nodes	2,703.0	1,304.4	378.8	95.8	91.0
	(2.1%)	(1.0%)	(0.29%)	(0.73%)	(0.69%)
z^G/z^*	1.0018	1.0017	1.0038	1.0049	1.0056
z^*/z_{LB}	1.07	1.05	1.025	1.0047	1.0007

problem. Where exactly comparable in the number of items and the number of periods, our branch-and-bound procedure is slightly faster. Upon final acceptance of this paper, we became aware of significant improvements in Raghavan and Rao's method, as reported in Raghavan and Rao (1992). The latter also describe a new solution approach based on a multi-commodity formulation of the problem using Dantzig-Wolfe decompositions. The latter method solves problems with 30 periods and 20 items in 480 seconds on the Sun 490 Sparc station. Moreover, this method is capable of solving problems with up to 40 periods and 100 items. (The latter sized problems require an average of approximately 7,500 seconds and involve linear programs with approximately 86,000 variables and constraints.) The partitioning heuristic described in the next section can be implemented with any exact solution method for the subproblems associated with each interval, i.e., the branch-and-bound method described in this section, or alternatively, the method of Raghavan and Rao (1992).

3. THE PARTITIONING HEURISTIC

In this section, we describe and analyze our proposed partitioning heuristic. The exact full or implicit enumeration procedures described in the previous section are only attractive if the length of the horizon is relatively small. This, together with our observation that forecast horizons tend to be small, suggests the following type of heuristic: Partition the complete horizon of N periods into relatively small intervals, and specify associated joint replenishment problems for each one of them. The lengths of these problems permit the use of the exact branch-and-bound method described in the previous section. The **JRP** associated with a given interval is obtained from the restriction of the complete **JRP** to the periods of that interval (see Figure 1). In addition, for each interval we add an option to choose starting conditions that appropriately complement the solutions obtained in prior intervals.

Indeed, we show below that an efficient partitioning heuristic of this type can be designed to be asymptotically optimal, or ε-optimal for any prespecified $\varepsilon > 0$. We also show that the partitioning heuristic compares favorably with the heuristics described in Section 2.
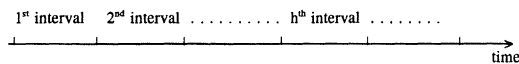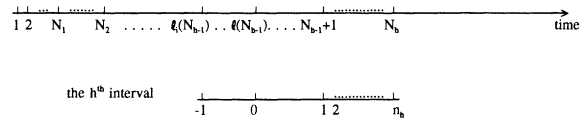

Figure 1. Time, partitioned into intervals.

Figure 2. Construction of the hth interval.

We now specify the proposed partitioning heuristic. We need the following auxiliary quantities:

$D_i(t) = \sum_{s=1}^{t} d_{is}$ = the cumulative demand for item i up to period t $(i = 1, \ldots, m; t = 1, \ldots, N)$;

$S_i(s, t) = \sum_{r=s}^{t-1} h_{ir} \sum_{k=r+1}^{t} d_{ik} = \sum_{r=s}^{t-1} h_{ir}(D_i(t) - D_i(r))$ = the total inventory carrying cost for item i under zero-inventory ordering in periods $s, \ldots, t$ when placing an order in period s to cover demands through period t $(s < t)$.

We partition the complete horizon $\{1, \ldots, N\}$ into I intervals of lengths $n_1, n_2, \ldots, n_I$ (i.e., $\sum_{h=1}^{I} n_h = N$) and specify associated joint replenishment problems denoted by $(\mathbf{JRP}_h)$. Let $N_h = \sum_{k=1}^{h} n_k$, $h = 1, \ldots, I$. We discuss possible choices for I, the number of intervals and their length, later in this section.

For $i = 1, \ldots, m$ let $\ell_i(N_{h-1})$ denote the last order period for item i in the partial solution constructed thus far, i.e., the solution constructed from $(\mathbf{JRP}_1)$ up to $(\mathbf{JRP}_{h-1})$.

Let $\ell(N_{h-1}) \equiv \max_i\{\ell_i(N_{h-1})\}$ denote the last of all order periods in the partial solution constructed thus far, and $L_{h-1} \equiv \{i: \ell_i(N_{h-1}) = \ell(N_{h-1})\}$ be the set of items ordered in that period. These quantities are clearly known after solving $\{(\mathbf{JRP}_{h-1})\}$ i.e., when specifying $(\mathbf{JRP}_h)$. $(\mathbf{JRP}_h)$, the joint replenishment problem specified for the hth interval, consists of $(n_h + 2)$ periods: The periods $N_{h-1} + 1, \ldots, N_h$ preceded by two dummy periods with zero demands and holding cost rates, which we refer to as periods -1 and 0 (see Figure 2).

Orders in period 0 represent additions to the orders placed in period $\ell(N_{h-1})$ to cover demands of some of the initial (or possibly all) periods in the hth interval. For items in the set L_{h-1} such additional orders can be made without *any* additional setup costs. For any product $i \notin L_{h-1}$, if additional orders are placed, only the item-specific setup cost $K_{i,\ell(N_{h-1})}$ needs to be incurred, while saving a *fixed* (not necessarily positive) amount of variable order and holding costs over the periods $\ell_i(N_{h-1}), \ldots, \ell(N_{h-1}) - 1$ for the units demanded in periods $\ell(N_{h-1}), \ldots, N_{h-1}$. The demands in these periods are, in the optimal solution for $(\mathbf{JRP}_{h-1})$, covered by an order in period $\ell_i(N_{h-1})$ and

can now be covered by a later order in period $\ell(N_{h-1})$. We thus specify the setup costs in period 0 as:

$K_{00} = 0$;

$K_{i0} = 0$ for $i \in L_{h-1}$;

$K_{i0} = K_{i\ell(N_{h-1})} + \kappa_{i0}$ for $i \notin L_{h-1}$;

where

$$\kappa_{i0} = (c_{i\ell(N_{h-1})} - c_{i\ell_i(N_{h-1})})$$
$$\cdot [D_i(N_{h-1}) - D_i(\ell(N_{h-1}) - 1)]$$
$$+ S_i(\ell_i(N_{h-1}), \ell(N_{h-1}) - 1)$$
$$+ S_i(\ell(N_{h-1}), N_{h-1})$$
$$- S_i(\ell_i(N_{h-1}), N_{h-1}). \qquad (9)$$

Observe that $K_{i0} \geq 0$ for $i \notin L_{h-1}$ because otherwise $\ell(N_{h-1})$ is a better last order period in $(\mathbf{JRP}_{h-1})$ than $\ell_i(N_{h-1})$.

The *variable* cost per unit of item i ordered in period 0 (for demands in the hth interval) consists of $c_{i\ell(N_{h-1})}$, the variable *order* cost of the corresponding period in the $(h - 1)$st interval, plus the cost of holding a unit over the interval $[\ell(N_{h-1}), N_{h-1}]$ i.e.,

$$c_{i0} = c_{i\ell(N_{h-1})} + h_{i\ell(N_{h-1})} + \cdots + h_{iN_{h-1}}.$$

Orders in period -1 represent additions to the orders for items $i \notin L_{h-1}$ in their respective last order periods $(<\ell(N_{h-1}))$ in the optimal solution of $(\mathbf{JRP}_{h-1})$. This period *only* considers additions to orders for items $i \notin L_{h-1}$ and these additions can be made with *no* (additional) setup costs involved. We thus specify the order cost parameters for period -1 as:

$K_{0,-1} = 0$;

$K_{i,-1} = \infty$ for $i \in L_{h-1}$;

$K_{i,-1} = 0$ for $i \notin L_{h-1}$;

$c_{i,-1} = c_{i\ell_i(N_{h-1})} + h_{i\ell_i(N_{h-1})} + \cdots + h_{i,N_{h-1}}$.

As mentioned, $d_{i0} = d_{i,-1} = h_{i,0} = h_{i,-1} = 0$ for all $i = 1, \ldots, m$.

After solving $(\mathbf{JRP}_h)$ *exactly* (e.g., via the branch-and-bound method described in Section 2) we update the partial solution for the horizon investigated so far. Updating the solution consists of two steps. In the first step we add the order quantities determined in $\mathbf{JRP}_h$ to the appropriate periods. Orders in periods $1, \ldots, n_h$ of the hth interval correspond with order quantities in periods $N_{h-1} + 1, \ldots, N_h$. For every item i, the order quantities for periods -1 and 0, if any, are added to the order quantities in periods $\ell_i(N_{h-1})$ and $\ell(N_{h-1})$, respectively. (Note that for

every item i, one of the order quantities in periods -1 or 0 must be zero, in view of the zero inventory ordering policy discussed in the Introduction and the fact that $d_{i,0} = 0$.)

The second step of updating the solution consists of determining the values of $\ell(N_h)$ and $\ell_i(N_h)$ for all i, according to the definitions of these quantities, with the understanding that for every i, periods -1 and 0 correspond with periods $\ell_i(N_{h-1})$ and $\ell(N_{h-1})$.

Finally, let $z(\mathbf{JRP}_h)$ denote the optimal cost for problem $(\mathbf{JRP}_h)$. We conclude:

Lemma 2. *The solution obtained by the partitioning heuristic is feasible and has a cost value $z^H \equiv \sum_{h=1}^{I} z(\mathbf{JRP}_h)$.*

We now derive worst-case bounds for the optimality gap which arises when I intervals are used in the partitioning heuristic. We do so under mild conditions with respect to the cost and demand parameters, specified in Assumptions 1 and 2.

Assumption 1. *There exist for all $i = 1, \ldots, m$ integers $M_i \geq 1$ and constants $K_{0^*}, K_{i^*}, d_{i^*}, h_{i^*}$ and c_{i^*} such that for all $t \geq 1$ and all $i = 1, \ldots, m$:*

$$(d_{it} + \cdots + d_{it+M_i}) \geq M_i d_{i^*} \text{ and } \sum_{t=1}^{N} d_{it} \geq N d_{i^*};$$

$$K_{0t} \geq K_{0^*}; \quad K_{it} \geq K_{i^*}; \quad h_{it} \geq h_{i^*}; \quad c_{it} \geq c_{i^*}.$$

Assumption 2. *There exist for all $i = 1, \ldots, m$ an integer $M_i \geq 1$ and constants $K_0^*, K_{0^*}, K_i^*, K_{i^*}, d_i^*, d_{i^*}, h_{i^*}, c_i^*$ and c_{i^*} such that for all $t \geq 1$ and all $i = 1, \ldots, m$:*

$$(d_{it} + \cdots + d_{it+M_i}) \geq M_i d_{i^*}; \quad d_{it} \leq d_i^*;$$

$$\sum_{t=1}^{N} d_{it} \geq N d_{i^*};$$

$$K_{0^*} \leq K_{0t} \leq K_0^*; \quad K_{i^*} \leq K_{it} \leq K_i^*;$$

$$c_{i^*} \leq c_{it} \leq c_i^*; \quad h_{i^*} \leq h_{it}.$$

Note that Assumption 1 is weaker than Assumption 2.

We first need to derive a lower bound for z^*.

Theorem 1. *(The lower bound theorem) Given Assumption 1,*

a. *If $M_i = 1$ for all $i = 1, \ldots, m$*

$$z^* \geq N\left(\max\left\{\sum_{i=1}^{m} \sqrt{2k_i h_{i^*} d_{i^*}} : \sum_{i=1}^{m} k_i = \sum_{i=0}^{m} K_{i^*}, \right.\right.$$
$$\left.\left. k_i \geq K_{i^*}\right\} + \sum_{i=1}^{m} c_{i^*} d_{i^*}\right) \equiv N\gamma_1.$$

b. If $M_i \geq 2$ for some $i = 1, \ldots, m$

$$z^* \geq N\left(\max\left\{ \sum_{i:M_i=1} \sqrt{2k_i h_{i^*} d_{i^*}} + \sum_{i:M_i \geq 2} \psi_i(k_i) \right.\right.$$

$$\left. : \sum_{i=1}^{m} k_i = \sum_{i=0}^{m} K_{i^*}, \; k_i \geq K_{i^*} \right\}$$

$$\left. + \sum_{i=1}^{m} c_{i^*} d_{i^*} \right) \equiv N\gamma_2,$$

where

$$\psi_i(k_i) = \begin{cases} k_i/2M_i & \text{if } K_i \leq k_i \leq M_i^2 h_{i^*} d_{i^*} \\ \sqrt{2(k_i + h_{i^*} M_i^2 d_{i^*})h_{i^*} d_{i^*}} - 1.5 h_{i^*} M_i d_{i^*} \\ & \text{if } k_i > M_i^2 h_{i^*} d_{i^*}. \end{cases}$$

Proof. a. A lower bound for z^* is obtained by sequentially implementing the following steps:

i. Replacing all cost and demand parameters by their corresponding lower bounds.

ii. Allocating the joint setup costs to the item-specific setup costs such that (see Section 1):

$$\sum_{i=1}^{m} k_i = \sum_{i=0}^{m} K_{i^*}$$

$$k_i \geq K_{i^*} \quad i = 1, \ldots, m. \tag{10}$$

iii. Transforming the problem into a continuous one, i.e., into a finite horizon EOQ problem, allowing for *continuous* rather than *discrete* review. It follows from Carr and Howe (1962) that this step represents a further lower bound for the optimal costs.

Finally, the lower bound stated in the theorem is obtained by choosing in step ii the vector k that maximizes the sum (over all items) of the minimum cost in each of the resulting m separate single item dynamic lot sizing models.

b. See the Appendix.

It is easy to verify that the functions $\psi_i(\cdot)$ are concave and continuously differentiable (by evaluating the left and right hand limits and derivatives in the breakpoint $k_i = M^2 h_{i^*} d_{i^*}$). Thus, to evaluate the bounds in parts a and b one needs to maximize a sum of concave and differentiable functions subject to a single budget constraint and lower bound constraints for the individual variables. The ranking method in Zipkin is ideally suited for this purpose; it requires $O(m\log m)$ elementary operations and at most m evaluations of square roots.

We are now ready to derive a worst-case bound for the heuristic's optimality gap. This bound serves, in addition, to suggest proper choices for the interval lengths to be employed in the heuristic; it is also the essential foundation for establishing that the heuristic can be designed to be asymptotically optimal, or ϵ-optimal for any $\epsilon > 0$. Without loss of generality we assume that a constant h_* exists such that $h_{it} \geq h_{i^*} = h_*$ for all $i = 1, \ldots, m$ and $t = 1, \ldots, N$, and that $d_{it} \geq 1$ or $d_{it} = 0$ for all $i = 1, \ldots, m$ and $t = 1, \ldots, N$. (This can be achieved by an appropriate choice of the units in which the items are measured. First, units are chosen so that all nonzero demands are ≥ 1; next, items i's unit is decreased by the factor $(\min_{1 \leq \ell \leq m} h_{\ell^*}/h_{i^*}) \leq 1$, which results in an increase of all demands by the reciprocal factor so that all nonzero demands continue to be ≥ 1.)

The worst-case bound and its proof are particularly simple if no speculative motive exists for carrying inventory for any of the items, i.e., in terms of variable costs, it is never advantageous to order some future period's demand in the current period, or $c_{it+1} \leq c_{it} + h_{it}$ for all $i = 1, \ldots, m$ and $t = 2, \ldots, N$. We therefore deal with this case first.

Theorem 2. *Let I denote the number of intervals employed by the partitioning heuristic. Assume that no speculative motives for carrying inventory prevail for any of the items. Assume that Assumption 1 is satisfied, and in addition there exist for all $i = 1, \ldots, m$ constants K_0^* and K_i^* such that for all $t \geq 1$, $K_{0t} \leq K_0^*$; and $K_{it} \leq K_i^*$. Then*

$$\frac{z^H - z^*}{z^*} \leq \frac{(I-1)}{N} \frac{\rho'}{\gamma},$$

where

$$\rho' = \sum_{i=0}^{m} K_i^* \text{ and } \gamma = \begin{cases} \gamma_1 & \text{if all } M_i = 1 \\ \gamma_2 & \text{otherwise} \end{cases}$$

(γ_1 and γ_2 defined as in Theorem 1).

Proof. We show that $(z^H - z^*) \leq (I - 1)\rho'$. The theorem then follows from Theorem 1. Consider an optimal solution of the **JRP** on the entire N-period horizon. If this optimal solution fails to be achievable by the partitioning heuristic, we can transform it into one that *is* achievable by this heuristic, adding no more than $\sum_{i=0}^{m} K_i^* = \rho'$ to the total cost. This can be achieved by adding a setup for each item in the first period of each of the $(I - 1)$ intervals (except the first one, in which the optimal solution is clearly achievable by the partitioning heuristic). Since in the transformed solution some of the orders are postponed, and no speculative motives are assumed, this

transformation does not add to the total *variable* order and/or holding costs.

The following theorem characterizes the worst-case upper bound for the case where speculative motives for carrying inventories may exist.

Theorem 3. *Let I denote the number of intervals employed by the partitioning heuristic, and assume that Assumption 2 is satisfied. Let* $\eta = \max_i\{c_i^* - c_{i*} - h_*\}$. *Then*

$$\frac{z^H - z^*}{z^*} \leqslant \frac{(I-1)}{N}\frac{\rho}{\gamma},$$

where

$$\rho = \sum_{i=0}^{m} K_i^* + \eta\left[L\sum_{i=1}^{m}d_i^* + \left(\sum_{i=0}^{m}K_i^*\right)\Big/h_*\right]$$

$$L = \lceil \max_i\{c_i^* - c_{i*}\}/h_*\rceil \text{ and } \gamma$$

$$= \begin{cases} \gamma_1 & \textit{if all } M_i = 1 \\ \gamma_2 & \textit{otherwise} \end{cases}$$

(γ_1 *and* γ_2 *defined as in Theorem* 1).

Proof. See the Appendix.

Remark. The bound for the optimality gap in Theorems 2 and 3 applies, in fact, to the more simplistic version of the partitioning heuristic in which no dummy periods (0 and −1) are added to any interval. Recall that these periods are added in the *proposed* version to allow for appropriate starting conditions. An alternative upper bound may be constructed by perturbing the optimal solution to the complete problem as follows: Instead of *postponing* the orders of any units carried into a given interval (say, interval h) to the first period of that interval (period $N_{h-1} + 1$), we may, for all items involved, *augment* the size of their last order as determined by the solution to the partitioning heuristic for intervals $1, \ldots, h-1$. This type of perturbation results in a solution achievable under the proposed version of the partitioning heuristic, and avoids *any* increase in setup costs. The increase in variable costs can be bounded by a technique similar to, though more tedious than, the one employed in the proof of Theorem 3. Moreover, the resulting bound is sometimes better, but sometimes worse, than the value of ρ in Theorem 3.

2.3. Possible Choices of Interval Lengths for the Partitioning Heuristic

Theorem 3 suggests the following choice for the interval lengths n_h ($h = 1, \ldots, I$) to be employed in the partitioning heuristic: (If no speculative motives

for carrying inventory exist, the arguments below are valid with ρ' replacing ρ, based on Theorem 2.)

$$n_h = \max\{Y, \lceil \log N\rceil\}, \quad h = 1, \cdots, I-1 \tag{11}$$

$$n_I = N - \sum_{h=1}^{I-1} n_h \tag{12}$$

with Y an arbitrary integer.

Theorem 3 shows that an ϵ-optimal solution may be guaranteed by choosing $Y = Y(\epsilon)$ where

$$Y(\epsilon) \equiv \min\{\lceil\rho/\epsilon\gamma\rceil, N\}. \tag{13}$$

Corollary 1. *Given Assumption* 2, *the partitioning heuristic results in an ϵ-optimal solution for any given $\epsilon \geqslant 0$ if the intervals n_h ($h = 1, \ldots, I$) are specified as in* (11) *and* (12) *and* $Y = Y(\epsilon)$.

Proof. Assume without loss of generality that $Y(\epsilon) = \lceil\rho/\epsilon\gamma\rceil$. If $Y(\epsilon) = N$, we clearly obtain an *optimal* solution. It follows from Theorem 3 that

$$\frac{z^H - z^*}{z^*} \leqslant \frac{(I-1)}{N}\frac{\rho}{\gamma} = \frac{(\lceil N/Y\rceil - 1)}{N}\frac{\rho}{\gamma}$$

$$\leqslant \frac{(N/Y)\rho}{N\gamma} \leqslant \frac{\rho}{Y\gamma} \leqslant \epsilon$$

because $Y \geqslant \rho/\epsilon\gamma$.

Theorem 3 also allows us to conclude that the partitioning heuristic is asymptotically optimal as N increases to infinity.

Corollary 2. *Consider the partitioning heuristic with interval lengths specified by* (11) *and* (12) *for any integer* Y.

a. *The heuristic has complexity* $O(mN^2\log\log N)$.
b. *Given Assumption* 2, *for fixed* m, *the heuristic is asymptotically optimal as* N *increases to infinity.*
c. *Given Assumption* 2 *with values* K_i^*, d_i^*, c_i^* *uniformly bounded in* i, *and* K_{i*}, d_{i*}, c_{i*} *uniformly bounded away from zero, the heuristic is asymptotically optimal as* $m \to \infty$ *and* $N \to \infty$.

Proof

a. The partitioning heuristic requires, to compute its solution for the hth interval, at most 2^{n_h} solutions of m single-item dynamic lot sizing models, each of which can be solved in $O(n_h\log n_h)$ time. The above complexity bound then follows because $n_h = O(\log N)$ and $I = O(N/\log N)$.
b. Asymptotic optimality for fixed m is immediate from Theorem 3.

c. Note that ρ is $O(m)$ while $\gamma = \Omega(m)$, i.e., there exists a constant $a > 0$ such that $\gamma \geq am$. Thus, $\rho/\gamma = O(1)$ as $m \to \infty$, and part c follows from part b.

The following example illustrates the magnitude of the upper bound for the worst-case optimality gaps associated with a given horizon length n, as determined by Corollary 1.

Example 1. Consider a problem instance with $m = 3$, $h_* = 1$, $M_i = 1$ $(i = 1, 2, 3)$, $K_{0*} = 70$, and $K_0^* = 80$. Bounds for the remaining parameters are specified as:

Item	K_i^*	K_{i*}	c_i^*	c_{i*}	d_i^*	d_{i*}
1	30	20	12	8	15	5
2	25	20	9	6	12	8
3	40	25	10	7	10	5

It is easy to verify (from the Kuhn-Tucker conditions) that $\gamma = 322.4$, while $\rho = 1,144$. This results in the following bounds for the optimality gaps, as a function of the chosen interval length n:

n	3	4	5	6	7	8	9	10	11	12	13
ϵ (%)	118	89	71	59	51	44	39	35	32	29	27

If no speculative motives for carrying inventories prevail, $\rho' = 175$ which results in these bounds for the optimality gap:

n	1	2	3	4	5	6	7	8	9	10
ϵ (%)	54	27	18	14	11	9	7.8	6.8	6	5.4

Observe that even for small interval lengths of $n = 5$ or $n = 10$ the bound for the optimality gap is no larger than 71% and 35%, respectively, when speculative motives prevail. When they do not, the bounds reduce to 11% and 5.4%, respectively.

Equations 11 and 12 describe a static procedure for the determination of the interval lengths to guarantee a specific optimality gap of the partitioning heuristic. These are, however, based on an a priori worst-case upper bound for $(z^H - z^*)$ and a worst-case a priori lower bound for z^*.

We conjecture that efficiency improvements may be achieved by specifying the interval lengths in a *dynamic* way. Dynamic specification of the interval lengths requires the upfront evaluation of the lower bound $z(\alpha^{*1}, \ldots, \alpha^{*N})$, as discussed in Section 1. Assume, therefore, that the vectors α^{*t} $(t = 1, \ldots, N)$ have been determined by solving (5) via the greedy procedure. Evaluation of the lower bound can now be performed by parallel executions for each of the m items of the forward single-item solution method of Federgruen and Tzur (1991). This permits keeping track, for $t = 1, \ldots, N$, of a lower bound $z^{LB}(t)$ for the minimum cost $z^*(t)$ in the **JRP** over the t-period horizon.

Assume now that the first N_0 periods have been partitioned into I_0 intervals, and that the partitioning heuristic has been applied to these, with an associated cost (for the first N_0 periods) $z^H(N_0)$. It follows from the proof of Theorem 3 that if the length of the next interval is n, then $z^H(N_0 + n) - z^*(N_0 + n) \leq z^H(N_0) - z^{LB}(N_0) + \rho$. Thus, if an ϵ-optimal overall solution is desired, the length of the next interval may be specified as

$$n^* = \min\left\{ n: \frac{z^H(N_0) - z^{LB}(N_0) + \rho}{z^{LB}(N_0 + n)} \leq \epsilon \right\}.$$

We have evaluated the performance of the partitioning heuristic with respect to a collection of 190 problem instances, grouped into 8 problem sets. These problem sets are generated from the same basic class of problems utilized in the numerical study of the branch-and-bound method in Section 2. The partitioning heuristic is applied twice or three times to each problem instance, for as many choices of the number of intervals I. We always choose intervals of equal length, the value of which is denoted by n in Tables V–XII.

In problem sets 1 and 2 (Tables V and VI) we vary the horizon length from $N = 18$ to $N = 24$ for problems with $m = 5$ and $m = 10$ items, respectively. In Table VII we return to the basic set (with $N = 18$ and $m = 5$) and vary α, the degree of intertemporal

Table V
Problem Set 1, $m = 5$

N	18		24			30		
	$n = 6$	$n = 9$	$n = 6$	$n = 8$	$n = 12$	$n = 6$	$n = 10$	$n = 15$
z^H/z^*	1.0041	1.00066	1.0058	1.0044	1.0016	1.0069	1.0028	1.0019
% opt.	40%	80%	0%	10%	50%	0%	30%	50%
CPU z^H	0.47	0.71	0.65	0.92	2.52	0.83	1.89	6.47
CPU z^*		4.0		40.1			231.7	

Table VI
Problem Set 2, $m = 10$

N	18		24			30		
	$n = 6$	$n = 9$	$n = 6$	$n = 8$	$n = 12$	$n = 6$	$n = 10$	$n = 15$
z^H/z^*	1.00082	1.0028	1.0017	1.0015	1.001	1.0015	1.0009	1.0011
% opt.	60%	30%	50%	50%	60%	60%	70%	70%
CPU z^H	0.81	1.84	1.18	1.9	12.74	1.46	3.94	18.24
CPU z^*		19.4		250.4			1112.75	

correlation between the cost and demand parameters, from 0 to 1. In problem sets 4–7 we vary the ranges of the demand and cost parameters (one at a time). Finally, in problem set 8 we consider the basic problem class with $m = 5$ items and *large* horizons, i.e., $N = 100$ and $N = 500$. Each problem set consists of 10 independently generated instances for each of the considered combinations of N, m, α and demand and cost parameters.

Like the branch-and-bound method, the partitioning heuristic has been encoded in FORTRAN and executed on an IBM 4381. Tables V–IX report on the ratio z^H/z^*, and the percentage of problem instances for which an optimal solution is found. The value of z^* is obtained from the numerical study of Section 2. In addition, for Tables V–VII we report the CPU time (in seconds) for the partitioning heuristic and the exact branch-and-bound method. In Table XII we report the ratio z^H/z_{LB} instead of z^H/z^* because the optimum value z^* is unavailable for the large horizon lengths considered there.

We are delighted to conclude that the partitioning heuristic performs exceptionally well, even when an interval length n of only six periods is employed. The average optimality gap (where measurable) is only 0.38% and in none of the problem categories is it larger than 0.78%. An optimal solution is found in 41.8% of the problem instances in this implementation of the partitioning heuristic. If intervals of length 9 or 10 are employed, the average optimality gap (where measurable) equals 0.23% and in none of the problem categories is it larger than 0.49% and in 52.67% of the instances is an optimal solution found.

There is even a problem category in which the partitioning heuristic obtains an optimal solution in all ten problem instances. This category has $\alpha = 0$ (see Table VII) and $n = 9$. As shown in Federgruen and Tzur (1994), forecast horizons are particularly small when α is small, i.e., when a high degree of intertemporal variability prevails in the cost and demand parameters. We believe this to be the explanation for the perfect performance of the partitioning heuristic in this problem category. Tables VII–XI exhibit a slight increase in the optimality gaps when the ranges of the parameters are increased.

For the last problem set the true optimality gap cannot be assessed, but its upper bound z^H/z_{LB} is still, on average, 3.3% and always smaller than 3.5%.

Table VII
Problem Set 3, $N = 18$, $m = 5$

α	0.0		0.2		0.5		0.8		1.0	
	$n = 6$	$n = 9$	$n = 6$	$n = 9$	$n = 6$	$n = 9$	$n = 6$	$n = 9$	$n = 6$	$n = 9$
z^H/z^*	1.0035	1.0	1.0054	1.0014	1.0041	1.00066	1.0007	1.0042	1.00006	1.0049
% opt.	60%	100%	40%	70%	40%	80%	80%	10%	90%	10%
CPU z^H	0.52	1.45	0.52	1.17	0.46	0.71	0.32	0.57	0.35	0.5
CPU z^*		25.1		13.2		4.0		1.2		1.1

Table VIII
Problem Set 4, $N = 18$, $m = 5$

K_i	(15, 25)		(10, 30)		(5, 35)	
	$n = 6$	$n = 9$	$n = 6$	$n = 9$	$n = 6$	$n = 9$
z^H/z^*	1.0034	1.0027	1.0041	1.00066	1.0063	1.0034
% opt.	50%	40%	40%	80%	30%	40%

Table IX
Problem Set 5, $N = 18$, $m = 5$

K_0	(90, 110)		(80, 120)		(70, 130)	
	$n = 6$	$n = 9$	$n = 6$	$n = 9$	$n = 6$	$n = 9$
z^H/z^*	1.0026	1.0019	1.0041	1.00066	1.0064	1.002
% opt.	50%	70%	40%	80%	20%	50%

Given our observation in Section 2 that the lower bound z_{LB} almost invariably comes within 2.4–2.8% of z^* for problems with $N \leqslant 36$, and the same value of α as the one used in problem set 8, we believe that the true average optimality gap for this problem set is no more than 0.7%.

As can be expected from the complexity analysis in this section and the numerical study in the previous one, the partitioning heuristic requires far less time than the branch-and-bound procedure. When intervals of length 6 are employed, the CPU time is always of the order of a single CPU second, even for problems with $N = 30$. Even when intervals of length $n = 9$ or 10 are used, the average CPU time is always less than 4 CPU seconds in all problem categories with $N \leqslant 30$ and no larger than 38.6 seconds for $N = 500$. As expected, larger intervals almost invariably result in smaller optimality gaps.

The overall desirability of the partitioning heuristic is perhaps best illustrated by considering the problem category with $N = 30$ and $m = 10$ in Table VI. An optimal solution requires more than a thousand CPU seconds, and with the partitioning heuristic one can get within 0.15% of optimality after approximately 1 CPU second only.

We conclude this section with some comparisons with alternative heuristics. Joneja (1990) implemented his heuristic (designed for the case of constant cost parameters) on an AT&T 6300 personal computer. On average, for his chosen set of problem instances, the heuristic has a solution that is 1.8% higher than a lower bound; the maximum deviation from the lower bound is 10.8%.

Likewise, the partitioning heuristic compares favorably with three alternative heuristics discussed in this paper: the greedy heuristic; a version of the branch-and-bound method of Section 2, designed to terminate as soon as a solution is found whose cost value comes within a prespecified relative gap ε of (a lower bound for) the optimum cost value; and the integer programming method of Raghavan and Rao (1992), again designed to terminate within an ϵ-optimal solution. As explained in Section 2, the greedy heuristic has complexity $O(mN^3\log N)$ while the complexity of the partitioning heuristic with interval lengths specified by (11) and (12) is $O(m N^2\log\log N)$ only (see Corollary 2). The partitioning heuristic is asymptotically optimal while the greedy heuristic is not. The former can be designed to generate an ϵ-optimal solution for any $\epsilon > 0$ while the greedy heuristic cannot, Proposition 1's interesting worst-case bound for the optimality gap of the greedy heuristic and its remarkable empirical performance, reported in Section 2, not withstanding. As mentioned in Section 2, the branch-and-bound method developed there, and the method of Raghavan and Rao (1992), can be used

Table X
Problem Set 6, $N = 18$, $m = 5$

c_i	(7, 8)		(5, 10)		(1, 14)	
	$n = 6$	$n = 9$	$n = 6$	$n = 9$	$n = 6$	$n = 9$
z^H/z^*	1.003	1.0017	1.0041	1.00066	1.0078	1.0016
% opt.	50%	30%	40%	80%	10%	80%

Table XI
Problem Set 7, $N = 18$, $m = 5$

d_i	(4, 7)		(1, 10)	
	$n = 6$	$n = 9$	$n = 6$	$n = 9$
z^H/z^*	1.0039	1.0031	1.0041	1.00066
% opt.	20%	40%	40%	80%

Table XII
Problem Set 8, $m = 5$

	100		500	
N	$n = 10$	$n = 20$	$n = 10$	$n = 20$
z^H/z_{LB}	1.033	1.03	1.035	1.034
CPU z^H	7.2	123.8	38.57	757.5

comfortably (as an exact method or ϵ-approximation scheme) for problems of up to moderate size, but their theoretical complexity is exponential, and their empirical complexity is several orders of magnitude larger than that of the partitioning heuristic.

APPENDIX

Proof of Theorem 1b

Before proving part b we need to derive a lower bound for the minimum cost in a *single*-item dynamic lot sizing model, when sporadic demand is allowed. To derive this lower bound we merely assume that the cumulative demand over a large enough time interval is uniformly bounded away from zero (see Lemma A1).

Lemma A1. *Consider a single-item dynamic lot sizing model over a horizon of N periods, and let z denote the minimum cost for this model. Let K_t, c_t, h_t, and d_t denote the setup cost, per unit order cost, holding cost rate and demand in period t ($t = 1, \ldots, N$). Assume that there exists an integer $M \geq 2$, and constants h_*, K_* c_* and d_* such that $(d_t + \cdots + d_{t+M}) \geq Md_*$, $\sum_{t=1}^{N} d_t \geq Nd_*$, $h_t \geq h_*$, $c_t \geq c_*$ and $K_t \geq K_*$ for all $t = 1, \ldots, N$. Then,*

$$z \geq \begin{cases} ((K_*/2M) + c_*d_*)N & \text{if } M > \sqrt{K_*/h_*d_*} \\ (\sqrt{2(K_* + h_*M^2d_*)h_*d_*} - 1.5h_*Md_* + c_*d_*)N & \\ & \text{if } M \leq \sqrt{K_*/h_*d_*}. \end{cases}$$

Moreover, if $M \leq \sqrt{K_*/h_*d_*}$, then

$$(\sqrt{2(K_* + h_*M^2d_*)h_*d_*} - 1.5h_*Md_*)N$$

$$\geq (\tfrac{1}{2} h_*d_*M)N.$$

Proof. We clearly obtain a lower bound by replacing all cost parameters by their corresponding lower bounds. We refer to the resulting model as the *transformed problem*. Consider a zero inventory ordering solution in which $I \geq 1$ orders are placed. For $\ell = 1, \ldots, I$ let n_ℓ denote the number of periods in the ℓth order cycle, i.e., the interval which contains the ℓth order period and all subsequent periods prior

to the next order period (if any). (The Ith interval terminates with period N.)

We first derive a lower bound for the total *holding* costs incurred in a single order cycle of n periods in the transformed problem. Renumber the periods in this cycle as $1, \ldots, n$ and let $n = \nu M + \ell$ with $0 \leq \ell < M$, i.e., $\nu = \lfloor n/M \rfloor$. Assume first that $\nu \geq 1$. Observe that in each interval $[(j - 1)M + \ell + 1, jM + \ell]$ ($j = 1, \ldots, \nu$) at least Md_* units are demanded. Since all those units are ordered in period 1, the lowest holding costs for these demands arise when all Md_* units are demanded in period $(j - 1)M + \ell + 1$ (i.e., in the *first* period of this interval) and none in the remaining periods of the interval $[(j - 1)M + \ell + 1, jM + \ell]$. It follows that the holding costs in a single-order cycle of n periods are bounded from below by

$$h_*(Md_*) \sum_{j=0}^{\nu-1} (\ell + jM)$$

$$= h_*Md_*[\nu\ell + \tfrac{1}{2}M\nu(\nu - 1)]$$

$$= h_*(Md_*)[\lfloor n/M \rfloor \ell + \tfrac{1}{2}M\lfloor n/M \rfloor(\lfloor n/M \rfloor - 1)]$$

$$\geq \tfrac{1}{2}h_*M^2d_*[(n/M - 1)^+(n/M - 2)^+]. \quad (A.1)$$

Clearly, the lower bound in (A.1) continues to hold when $\nu = 0$.

This lower bound for the holding costs in a single order cycle implies the following lower bound for the *total* cost over the complete horizon, associated with the above described zero inventory ordering solution:

$$K_*I + \sum_{\ell=1}^{I} \tfrac{1}{2}h_*M^2d_*(n_\ell/M - 1)^+(n_\ell/M - 2)^+$$

$$+ c_*d_*N. \quad (A.2)$$

(The last term in (A.2) represents a lower bound for the variable order costs.) We conclude that

$$z \geq c_*d_*N$$

$$+ \min_I \left\{ K_*I + \min \left\{ \sum_{\ell=1}^{I} g(n_\ell) : \sum_{\ell=1}^{I} n_\ell \right. \right.$$

$$\left. \left. = N; n_\ell \text{ integers} \right\} \right\},$$

where $g(x) \equiv \frac{1}{2}h_*M^2d_*(x/M - 1)^+(x/M - 2)^+$ is convex.

A further lower bound is obtained by relaxing the integrality constraints for the n_ℓ-variables. It is well known (in view of the convexity of $g(\cdot)$) that $\sum_{\ell=1}^I g(n_\ell)$ is minimized by setting all $n_\ell = N/I(\ell = 1, \ldots, I)$. We thus obtain:

$$z \geq c_*d_*N + \min_{I \geq 1}$$

$$\cdot \{K_*I + (\frac{1}{2}h_*M^2d_*)I(N/IM - 1)^+(N/IM - 2)^+\}.$$

$$(A.3)$$

Note that the minimum in (A.3) cannot be achieved for $I \geq N/2M$ because for such values of I, the term K_*I can be decreased without increasing the other terms in I. Thus,

$$z \geq c_*d_*N + \min_{I \leq N/2M}$$

$$\cdot \{K_*I + (\frac{1}{2}h_*M^2d_*)[N^2/M^2I - 3N/M + 2I]\}$$

$$= c_*d_*N + \min_{I \leq N/2M}$$

$$\cdot \{\hat{K}I + (\hat{H}M^2)[N^2/M^2I - 3N/M]\}, \qquad (A.4)$$

where $\hat{K} = K_* + h_*M^2d_*$ and $\hat{H} = \frac{1}{2}h_*d_*$.

The minimum in (A.4) is clearly achieved for

$$I = N \min \{(\hat{H}/\hat{K})^{1/2}, 1/(2M)\}.$$

Substituting this value of I into (A.4) we obtain after some algebra,

$$z \geq \begin{cases} (K_*/2M)N + c_*d_*N & \text{if } (\hat{H}/\hat{K})^{1/2} \geq 1/(2M) \\ (2\sqrt{\hat{K}\hat{H}} - 3\hat{H}M + c_*d_*)N \geq (\hat{H}M + c_*d_*)N \\ & \text{if } (\hat{H}/\hat{K})^{1/2} < 1/(2M). \end{cases}$$

$$(A.5)$$

(To verify the inequality in (A.5), note that

$$2\sqrt{\hat{K}\hat{H}} - 3\hat{H}M = \hat{H}(2\sqrt{\hat{K}/\hat{H}} - 3M) > \hat{H}(4M - 3M)$$

$$= \hat{H}M.$$

The proof of Theorem 1b now follows from allocating the joint setup costs to item-specific setup costs as in (10), applying Lemma A1 to each of the m resulting separate single-item dynamic log sizing models, and taking the vector k in (10) that maximizes this bound.

Proof of Theorem 3

We show that $(z^H - z^*) \leq (I - 1)\rho$. The theorem then follows from Theorem 1. Consider an optimal solution of the **JRP** on the entire N-period horizon. We show that this solution can be transformed into

one which is achievable by the partitioning heuristic, adding at most $(I - 1)\rho$ to the total cost.

If the optimal solution to the complete **JRP** fails to be achievable by the partitioning heuristic, there must be order cycles for some or all items which start at some interval, say interval h $(h = 1, \ldots, I)$, and terminate in a later one. To transform the solution into one which is achievable by the partitioning heuristic, add an order in the first period of the $(h + 1)$st interval for all such items (with a crossing order cycle) (see Figure 3). This transformation adds at most $\sum_{i=0}^m K_i^*$ to the total setup costs. In case speculative motives prevail for some of these items, the variable cost may increase as well, but only for units demanded in the $(h + 1)$st interval because the units demanded in periods belonging to the hth interval continue to be ordered in the same period as before the transformation.

Let $S = \{i_1, \ldots, i_{|S|}\}$ denote the set of items with a crossing order cycle, and let t_ℓ denote the number of periods in the $(h + 1)$st interval which are included in the crossing order cycle of item i_ℓ (see Figure 3). Assume the items are numbered such that $t_{i_1} \leq t_{i_2} \leq \cdots \leq t_{i_{|S|}}$.

Let D^{i_ℓ} $(\ell = 1, \ldots, |S|)$ denote the number of units of item i_ℓ carried in inventory at the beginning of the $(h + 1)$st interval. Renumber the periods in the $(h + 1)$st interval from $1, \ldots, n_{h+1}$ and let $S_\ell = \{i_\ell, \ldots, i_{|S|}\}$. Note that for all $1 \leq r \leq t_\ell$:

$$rh_* \sum_{i \in S_\ell} (D^i - X_r^i)$$

$$\leq \sum_{i \in S_\ell} K_i^* + K_0^* + \sum_{i \in S_\ell} (c_i^* - c_{i^*})(D^i - X_r^i)$$

$$\leq \sum_{i \in S_\ell} K_i^* + K_0^* + Lh_* \sum_{i \in S_\ell} (D^i - X_r^i), \qquad (A.6)$$

where X_r^i denotes the cumulative demand for item i in the first $(r - 1)$ periods of the $(h + 1)$st interval. (If the first inequality is violated, then a strict cost improvement can be achieved by placing an order for all items in S_ℓ in period r, thus cutting each crossing order cycle for these items into two. Note that the

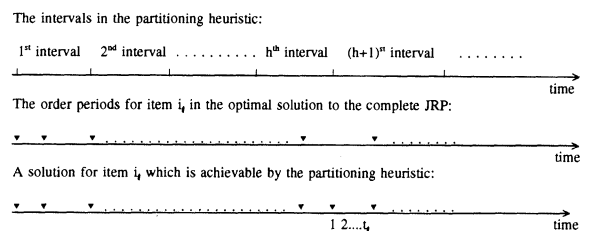The intervals in the partitioning heuristic:

Figure 3. Transformation of the optimal solution.

additional setup costs incurred due to this transformation are bounded by $\sum_{i \in S_\ell} K_i^* + K_0^*$. The additional variable costs are bounded by $\sum_{i \in S_\ell} (c_i^* - c_{i^*})(D^i - X_r^i)$ because exactly $(D^i - X_r^i)$ units of item i are now ordered in period r rather than at period b_ℓ, the beginning of the crossing order cycle. The transformation reduces the inventory level of item i by $(D^i - X_r^i)$ in each period from b_ℓ until $(r - 1)$, i.e., in at least r periods, resulting in a cost saving of at least $rh_*(D^i - X_r^i)$.)

Let $\bar{X}_r^i = D^i - X_r^i$. It follows from (A.6) that

$$(r - L)H_* \sum_{i \in S_\ell} \bar{X}_r^i \leq \sum_{i \in S_\ell} K_i^* + K_0^*. \tag{A.7}$$

Thus, for $r > L$,

$$\sum_{i \in S_\ell} \bar{X}_r^i \leq \frac{\sum_{i \in S_\ell} K_i^* + K_0^*}{(r - L)h_*}, \tag{A.8}$$

and

$$\sum_{i \in S_\ell} X_r^i \leq (r - 1) \sum_{i \in S_\ell} d_i^*. \tag{A.9}$$

Adding (A.8) and (A.9) and taking the minimum over r in the interval $[L + 1, t_\ell]$ we get:

$$\sum_{i \in S_\ell} D^i \leq \min_{L+1 \leq r \leq t_\ell} \left\{ (r - 1)\left(\sum_{i \in S_\ell} d_i^* \right) + \frac{\sum_{i \in S_\ell} K_i^* + K_0^*}{(r - L)h_*} \right\}. \tag{A.10}$$

The expression within curled brackets is clearly a convex function of r which decreases for

$$L + 1 \leq r \leq r^* \equiv L + \left[\sqrt{\frac{2\left(\sum_{i \in S_\ell} K_i^* + K_0^* \right)}{(h_* \sum_{i \in S_\ell} d_i^*)}} \right]$$

or

$$r^* \equiv L + \left[\sqrt{\frac{2\left(\sum_{i \in S_\ell} K_i^* + K_0^* \right)}{(h_* \sum_{i \in S_\ell} d_i^*)}} \right].$$

Assume first that $t_\ell \geq L + 1$. It follows that the minimum in (A.10) is achieved for $r = \min(t_\ell, r^*)$, i.e.,

$$\sum_{i \in S_\ell} D^i \leq [\min(t_\ell, r^*) - 1]\left(\sum_{i \in S_\ell} d_i^* \right)$$

$$+ \frac{\sum_{i \in S_\ell} K_i^* + K_0^*}{(\min(t_\ell, r^*) - L)h_*} \quad t_\ell \geq L + 1. \tag{A.11}$$

If $t_\ell \leq L$ we clearly have the bound

$$\sum_{i \in S_\ell} D^i \leq t_\ell \sum_{i \in S_\ell} d_i^*. \tag{A.12}$$

Let $b_\ell(t_\ell)$ denote the right-hand side of (A.11) when $t_\ell \geq L + 1$ and of (A.12) when $t_\ell \leq L$. For each demand unit of item $i \in S$ whose procurement costs are increased due to the transformation, the increase is clearly bounded by $(c_i^* - c_{i^*} - h_*)$. A bound for the increase in variable costs due to the transformation is thus given by the value of the linear program:

$$\Phi(t_1, \ldots, t_{|S|}) \equiv \max \sum_{i \in S} (c_i^* - c_{i^* - h_*})D^i$$

subject to

$$\sum_{i \in S_\ell} D^i \leq b_\ell(t_\ell), \quad \ell = 1, \ldots, |S| \tag{A.13}$$

$$D^i \geq 0.$$

Note that the function $b_\ell(\cdot)$ achieves its maximum for $t_\ell = L + 1$, because it is increasing for $t_\ell \leq L$, see (A.12), decreasing for $t_\ell \geq L + 1$ and

$$b_\ell(L + 1) = L \sum_{i \in S_\ell} d_i^* + \left(\sum_{i \in S_\ell} K_i^* + K_0^* \right) \Big/ h_*$$

$$> L \sum_{i \in S_\ell} d_i^* = b_\ell(L).$$

The optimum value of the linear program is clearly nondecreasing in the values $\{b_\ell: \ell = 1, \ldots, |S|\}$. It follows that $\Phi(L + 1, L + 1, \ldots, L + 1) = \max\{\Phi(t_1, \ldots, t_{|S|}) : t_1 \leq t_2 \leq \cdots \leq t_{|S|}\}$ because the function Φ achieves its unconstrained maximum in the point $(L + 1, \ldots, L + 1)$. Thus, substituting $t_\ell = L + 1$ in (A.13) we obtain the following bound for the total increase in variable costs due to the transformation:

$$\max \sum_{i \in S} (c_i^* - c_{i^*} - h_*)D^i$$

subject to

$$\sum_{i \in S_\ell} D^i \leq L \sum_{i \in S_\ell} d_i^*$$

$$+ \left(\sum_{i \in S_\ell} K_i^* + K_0^* \right)/h_* \quad \ell = 1, \ldots, |S|.$$

$$D^i \geq 0. \tag{A.14}$$

The objective function of (A.14) can clearly be bounded from above by $\eta \sum_{i \in S} D^i$. The latter objective is clearly maximized when

$$\sum_{i \in S} D^i = L \sum_{i \in S} d_i^* + \left(\sum_{i \in S} K_i^* + K_0^* \right) \Big/ h_*.$$

Thus, the resulting bound is clearly maximized when $S = \{1, \ldots, m\}$ in which case it equals to

$$\eta \left\{ L \sum_{i \in S} d_i^* + \left(\sum_{i \in S} K_i^* + K_0^* \right) \Big/ h_* \right\}.$$

(The optimum value of (A.14) itself can be obtained in closed form because the polyhedron described by its constraints is a polymatroid (see Edmonds 1970 or Frank and Tardos 1988). The linear program (A.14) thus can be solved by the greedy procedure. The resulting optimum value, when maximized over all permutations $(i_1, \ldots, i_m)$ results in the same bound.)

ACKNOWLEDGMENT

REFERENCES

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

AKSOY, Y., AND S. S. ERENGUC. 1988. Multi-Item Inventory Models With Coordinated Replenishments: A Survey. *Int. J. Opns. and Prod. Mgmt.* **8**, 63–73.

ARKIN, E., D. JONEJA AND R. ROUNDY. 1989. Computational Complexity of Uncapacitated Multi-Echelon Production Planning Problems. *O. R. Letts.* **8**, 61–66.

ATKINS, D. R., AND P. O. IYOGUN. 1988. A Heuristic With Lower Bound Performance Guarantee for the Multiproduct Dynamic Lot-Size Problem. *IIE Trans.* **20**, 369–373.

BROWN, R. G. 1967. *Decision Rules for Inventory Management.* Holt, Rinehart and Winston, New York, 46–55.

CARR, C. R., AND C. W. HOWE. 1962. Optimal Service Policies and Finite Time Horizons. *Mgmt. Sci.* **9**, 126–140.

CORNUEJOLS, G., M. L. FISHER AND G. L. NEMHAUSER. 1977. Location of Bank Accounts to Optimize Float: An Analytic Study of Exact and Approximate Algorithms. *Mgmt. Sci.* **23**, 789–810.

EDMONDS, J. 1970. Submodular Functions, Matroids and Certain Polyhedra. In *Combinatorial Structures and Their Applications*, R. Guy et al. (eds.). Gordon and Breach, New York, 69–87.

ERENGUC, S. S. 1988. Multiproduct Dynamic Lot-Sizing Model With Coordinated Replenishments. *Naval Res. Logist.* **35**, 1–22.

FEDERGRUEN, A., AND H. GROENEVELT. 1986. The Greedy Procedure of Resource Allocation Problems: Necessary and Sufficient Conditions for Optimality. *Opns. Res.* **34**, 909–919.

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.

FEDERGRUEN, A., AND M. TZUR. 1994. Minimal Forecast Horizons and a New Planning Procedure for the General Dynamic Lot Sizing Model: Nervousness Revisited. *Opns. Res.* **42**, 456–468.

FEDERGRUEN, A., AND Y. S. ZHENG. 1992. The Joint Replenishment Problem With General Joint Cost Structures. *Opns. Res.* **40**, 384–403.

FOX, B. 1988. Discrete Optimization via Marginal Analysis. *Mgmt. Sci.* **4**, 136–153.

FRANK, A., AND E. TARDOS. 1988. Generalized Polymatroid and Submodular Flows. *Math. Prog.* **42**, 489–563.

FREDERICKSON, G. N., AND D. B. JOHNSON. 1982. The Complexity of Selection and Ranking in $X + Y$ and Matrices With Sorted Columns. *J. Comput. and Sys. Sci.* 197–208.

GALIL, Z., AND N. MEGIDDO. 1979. A Fast Selection Algorithm and the Problem of Optimum Distribution of Efforts. *J. ACM* **26**, 58–64.

GROENEVELT, H. 1986. Resource Allocation Problems With Decreasing Marginal Returns to Scale. Ph.D. Dissertation, Columbia University, New York.

GROSS, O. 1956. A Class of Discrete-Type Minimization Problems. Working Paper RM-1644-PR, The Rand Corporation, Santa Monica, California.

IYOGUN, P. 1987. Lower Bounds for Production/ Inventory Problems by Cost Allocation. Ph.D. Dissertation, University of British Columbia, Vancouver, BC, Canada.

JACKSON, P. L., W. L. MAXWELL AND J. A. MUCKSTADT. 1985. The Joint Replenishment Problem With a Powers of Two Restriction. *IIE Trans.* **17**, 25–32.

JONEJA, D. 1989. Planning for Joint Replenishment and Assembly Systems With Deterministic Nonstationary Demands. Ph.D. Dissertation, Cornell University, Ithaca, New York.

JONEJA, D. 1990. The Joint Replenishment Problem: New Heuristics and Worst Case Performance Bounds. *Opns. Res.* **38**, 711–723.

KAO, E. P. C. 1979. A Multiproduct Dynamic Lot-Size Model With Individual and Joint Setup Costs. *Opns. Res.* **27**, 279–289.

LEOPOULOS, V. J., AND J. M. PROTH. 1985. The General Multiproducts Dynamic Lot Size Model With Individual Inventory Costs and Joint Productions Costs. *RAIRO APII* **19**, 117–130.

NEMHAUSER, G. L., AND L. A. WOLSEY. 1988. *Integer and Combinatorial Optimization*. John Wiley, New York.

NEMHAUSER, G. L., L. A. WOLSEY AND M. L. FISHER. 1978. An Analysis of Approximations for Maximizing Submodular Set Functions—I. *Math. Prog.* **14**, 265–294.

RAGHAVAN, P., AND M. R. RAO. 1991. The Multi-Item Lot Sizing Problem With Joint Replenishments: A Polyhedral Approach. Technical Report SOR-91-8, Department of Statistics and Operations Research, Leonard N. Stern School of Business, New York University, New York.

RAGHAVAN, P., AND M. R. RAO. 1992. Formulations for Multi-Item Lot-Sizing Problem With Joint Replenishment. Working Paper SOR-92-19, Stern School of Business, New York University, New York.

ROUNDY, R. 1985. 98% Effective Integer Ratio Lot Sizing for One Warehouse Multiretailer Systems. *Mgmt. Sci.* **31**, 1416–1430.

SILVER, E. A. 1976. Some Thoughts on Extensions of Lot Sizing Procedures Under Deterministic Time Varying Demand. Working Paper No. 104, Department of Management Science, University of Waterloo, Canada.

SILVER, E. A. 1979. A Simple Inventory Replenishment Decision Rule for a Linear Trend in Demand. *J. Opns. Res. Soc.* **30**, 71–75.

TZUR, M. 1992. A New Planning Approach for Single and Multi-Item Dynamic Lot-Sizing Models. Ph.D. Dissertation, Columbia University, New York.

VEINOTT, J. R. 1969. Minimum Concave-Cost Solution of Leontief Substitution Models of Multifacility Inventory Systems. *Opns. Res.* **17**, 262–291.

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.

ZANGWILL, W. I. 1966. A Deterministic Multiproduct Multifacility Production and Inventory Model. *Opns. Res.* **14**, 486–507.

ZIPKIN, P. H. 1980. Simple Ranking Methods for Allocation of One Resource. *Mgmt. Sci.* **26**, 34–43.