# Approximate solutions for a stochastic lot-sizing problem with partial customer-order information

**Approximate solutions for a stochastic lot-sizing problem with partial customer-order information**

N.P. Dellaert, M.T. Melo

WP 38



| | |
|---|---|
| BETA publicatie | WP 38 (working paper) |
| ISBN | 90-386-1009-2 |
| ISSN | 1386-9213 |
| NUGI | 684 |
| Eindhoven | Oktober 1999 |
| Keywords | - |
| BETA-Research Programme | Unit Management |
| Te publiceren in: | Operations Research |

# Approximate solutions for a stochastic lot-sizing problem with partial customer-order information

N.P. Dellaert

Eindhoven University of Technology

P.O. Box 513, 5600 MB Eindhoven, The Netherlands

M.T. Melo

Institute for Techno- and Economathematics (ITWM)

Erwin-Schroedinger-Strasse, Building 49

D-67663 Kaiserslautern, Germany

**Abstract**

We address a stochastic single item production system in a make-to-stock environment with partial knowledge on future demand resulting from customers ordering in advance of their actual needs. The problem consists of determining the optimal size of a production lot to replenish inventory, so that delivery promises are met on time at the expense of minimal average costs. For this problem an optimal policy is formulated. However, since the optimal policy is likely to be too complex in most practical situations, we present approximate strategies for obtaining production lot sizes. The well-known $(R, S)$ inventory policy is compared to two rules where production decisions take into account the available information on future customer requirements and the probabilistic characterisation of orders yet to be placed. It is shown that the $(s, S)$ inventory policy is a special case of one of the rules. An extensive numerical study reveals that the newly developed strategies outperform the classical ones.

# 1 Introduction

We consider the production of an item on a single machine with unrestricted capacity. Production activity is carried out in discrete time periods. Customer orders arrive randomly and are divided into different categories according to their degree of urgency. If the lead times that are promised by the company vary from 1 to $N$ periods, then the customer orders can be divided into $N$ categories. In an arbitrary period $t$, the company promises to customers of category $i$ $(1 \leq i \leq N)$ to have their orders ready by the end of period $t+i$. Furthermore, demands in each order category are assumed to be stationary, independently distributed and specified by known probability distributions. Demand categories can arise due to market strategy reasons, as in those situations where customers receive a special treatment by paying a premium or by signing long-term purchase contracts. In many situations, customers have a good perception of their needs and can provide advance warning of their actual demands. Gregory et al. [7] describe such a situation in a knitwear company. Furthermore, with the advent of electronic data interchange and private satellite communication systems, it became possible to link the computers of customers and suppliers so that orders can be placed instantaneously, thus providing suppliers with additional advance warning. Also, in JIT manufacturing, purchase contracts typically specify exact intervals between orders and deliveries. An example of a company where customer orders are divided into categories with fixed delivery promises is given by Dellaert & Wessels [5] who report the results of a simulation study for a lot-sizing problem in make-to-order manufacturing with multiple items and machines.

At the end of each period the system is reviewed. If the decision is to operate the machine in the following period, then a certain batch of items is manufactured. The items can be used for immediate delivery or alternatively can be added to a buffer stock of finished products. This buffer stock serves to anticipate unknown short term demand. In every period with positive production a fixed setup cost is charged which accounts for preparation costs associated with machine adjustments, cleaning, etc. Customer orders that are not ready by their due dates are backlogged and subject to a linear penalty cost until delivery. In addition, items kept in stock pay linear holding costs. The problem concerns the amount of items to be produced, given the pattern of known customer orders and a probabilistic characterization of orders yet to come. We seek a lot-sizing policy that minimizes the long-run average costs per period.

An interesting feature of our problem is that it contains partial knowledge on future customer orders, unlike the pure make-to-stock situation where only statistical information about future demand is available or used. At the end of an arbitrary period $t$, we know the exact demand requirements for period $t+1$ and also have partial knowledge on the required deliveries for periods $t+2,\dots,t+N$. The incomplete information about the demand derives from the possibility of receiving more orders for these periods, after $t$. Dellaert & Melo [3, 4] explored this characteristic in a make-to-order context. Unlike pure inventory models, it seems natural that our production decisions should take into account the pattern of known customer orders.

Recently, Hariharan & Zipkin [9] and Wijngaard [18] investigated the effect of foreknowledge of demand in make-to-stock environments. The problems studied by Hariharan & Zipkin [9] rely on the assumption that customers are all equally important which is a special case of our model. Hariharan & Zipkin analyze the effect of demand lead time on overall system performance for both the cases of constant and independent stochastic lead times. Wijngaard [18] also assumes the same demand lead time for all orders and analyses the role of foreknowledge of demand in situations with high utilization rate. One of the benefits of knowing future customer requirements is that it gives the possibility to anticipate capacity problems by initiating production of known orders in advance of their due dates. Zijm [19] investigated a stochastic lot-sizing problem with the same demand categories as in our model. By assuming that in each period with positive production a fixed amount of orders is manufactured, Zijm obtains myopic decisions by means of dynamic programming. Although this technique appears to give near-optimal results, it does not work well in our case since we charge a setup cost in every period with production. Recent studies based upon the same dynamic demand evolution have been carried out by Heath and Jackson [11], Güllü [8] and Graves, Kletter and Hetzel [6]. However, they all use demand forecasts and forecast revisions, instead of *known* probability distributions. Another difference between these three papers and our paper is the fact that we assume that the demand for the first period is always known completely, whereas they assume that this is not the case. The consequence of this difference will be discussed in Section 5. Heath and Jackson considered a dynamic forecast process as part of a simulation model that was used to analyze safety stock levels in a multiproduct situation and Güllü studied the value of extra information in the forecast evolution. Graves et al. analyzed a single-stage model to determine production smoothness and stability. Afterwards they show that the single stage can be used as a building block in a multistage supply chain.

In the next section, we formulate our problem as a Markov Decision Process with multi-dimensional state space, and briefly sketch how the optimal production/inventory policy can be obtained. Since the optimal policy is likely to be too complex in most practical situations, we present in Section 3 three different lot-sizing strategies yielding approximate solutions for our problem. The first strategy is the so-called $(x, T, w)$-rule where production is initiated during a period for which the required deliveries are at least $x$ units. In that situation, the known demand for the next $T$ periods is manufactured plus $w$ units for a buffer stock. In case $T = 1$, the $(x, T, w)$-rule reduces to the well-known $(s, S)$ policy. The second rule is a least-cost-per-period strategy which allows production to cover the known demand requirements for a non-fixed number of periods. Furthermore, unlike the $(x, T, w)$-rule, the amount manufactured for the buffer stock is not always the same. In our comparison we also include the well-known $(R, S)$ inventory policy, which neglects all future demand information. We first selected some small test problems to compare the performance of the heuristics to the optimal solution. Then we tested the lot-sizing heuristics by selecting a larger data set covering a wide variety of demand and cost parameters. The results obtained are reported in Section 4. Finally, Section 5 summarizes the main conclusions.

# 2 The Markov decision model and the optimal policy

We denote by $D_1, D_2, \ldots, D_N$ the one-period demands for the item in each order category, and assume that they are independent, non-negative discrete random variables with known probability distributions such that $d_i(\ell) = \mathbb{P}(D_i = \ell)$ represents the probability that during an arbitrary period $t$, $\ell$ orders belonging to category $i$ are placed to be delivered by the end of period $t+i$, with $1 \leq i \leq N$, $\ell = 0, 1, \ldots$ and $\sum_{\ell \geq 0} d_i(\ell) = 1$. For simplicity, we consider that each customer order consists of one unit of product. Observe that $d_i(\ell)$ does not depend on the time period $t$ due to the stationarity assumption. Therefore, we shall suppress time subscripts throughout. Within each period the following sequence of events occurs:

1) implementation of lot-sizing decision;
2) completion of production (if any);
3) inventory replenishment (if production took place);
4) arrival of new demand;
5) review of the system and cost accounting.

We formulate our problem as a dynamic program with stages being successive time periods and states being described by $N$-dimensional vectors $\boldsymbol{r} = (r_1, r_2, \ldots, r_N)$. At the end of an arbitrary period $t$, the first component $r_1$ is defined by

$$r_1 := \text{required deliveries for period } (t+1) \ + \text{backorders} \ - \text{stock}.$$

A negative value of $r_1$ should be interpreted as the amount of items held in stock. All the other components $r_i$ of the order state vector give the number of known but unfilled customer orders with a due date $i$ periods ahead, that is, $t+i$ with $2 \leq i \leq N$.

Associated with each state $\boldsymbol{r}$ there is a non-empty set of actions $A(\boldsymbol{r})$. An action $a$ specifies the size of a production lot. In case the inventory level in a certain period is non-negative (that is, $-r_1 \geq 0$), this means that there is no demand for immediate delivery. In such a situation the best action is to delay production. This results from the fact that production capacity is unrestricted, from our exact knowledge of the demand requirements for the next period, and from the assumption that when production occurs in a period, the finished items are available before the end of the same period. Hence, action $a = 0$ is to be selected when $r_1 <= 0$. If $r_1 > 0$, we can exclude action $a = 0$ in those situations where the penalty costs charged for not satisfying this demand on time are higher than the setup cost. Denoting by $K$ the fixed setup cost and by $p$ the unit penalty cost charged to each backlogged item, it follows that the action space $A(\boldsymbol{r})$ for each state $\boldsymbol{r}$ is defined as:

$$\begin{array}{ll} A(\boldsymbol{r}) = \{0\} & \text{if } r_1 \leq 0, \\ A(\boldsymbol{r}) = \{r_1, r_1+1, \ldots, \sum_{i=1}^{N} r_i, \ldots\} & \text{if } r_1 p > K, \\ A(\boldsymbol{r}) = \{0, r_1, r_1+1, \ldots, \sum_{i=1}^{N} r_i, \ldots\} & \text{if } r_1 p \leq K. \end{array}$$

Although theoretically $A(\boldsymbol{r})$ may contain an infinite number of actions, the knowledge of the demand distributions provides in practice information on the subset of actions that actually needs to be considered.

Moving ahead one period causes production to be netted against requirements. Demand for the previous period that was not satisfied is carried forward as backlogged demand in the current period. Let $\boldsymbol{Q}_a(\boldsymbol{r})$ denote the order state before the arrival of new demand. In the sequence of events listed above, $\boldsymbol{Q}_a(\boldsymbol{r})$ results from applying 3). At the end of every period, new demands $(\ell_1, \ell_2, \ldots, \ell_N)$ for each one of the next $N$ periods are placed. These are added to any unfilled demands observed in previous periods. Thus,

$\boldsymbol{Q}_a(\boldsymbol{r}) = (r_1 - a + r_2, r_3, \ldots, r_N, 0)$ and the transition to a new state $\boldsymbol{z} = \boldsymbol{Q}_a(\boldsymbol{r}) + (\ell_1, \ldots, \ell_N)$ takes place with probability $P^a_{\boldsymbol{rz}} = \prod_{i=1}^N d_i(\ell_i)$.

The direct costs $C^a_{\boldsymbol{r}}$ of selecting action $a$ on observing a state $\boldsymbol{r}$ are charged in the period in which they are incurred as in a classical inventory model. If $a = 0$ and $r_1 \leq 0$, we have holding costs $-hr_1$ with $h$ denoting the unit holding cost. In case $a = 0$ but $r_1 > 0$, we pay penalty costs $pr_1$. Finally, manufacturing at least $r_1$ items leads to a setup cost and holding costs for those $a - r_1$ items that will replenish the inventory. In summary, we have

$$C^a_{\boldsymbol{r}} = \begin{cases} -hr_1 & \text{if } a = 0 \text{ and } r_1 \leq 0, \\ pr_1 & \text{if } a = 0 \text{ and } r_1 > 0, \\ K + h(a - r_1) & \text{if } a \geq r_1. \end{cases}$$

To determine the lot-sizing policy with minimal long-run average costs, we apply the value-iteration algorithm as described by Odoni [13] and Tijms [16, Chapter 3]. Let $v_n(\boldsymbol{r})$ denote the minimal expected costs with $n$ periods left to the time horizon when the current state is $\boldsymbol{r}$ and a terminal cost $v_0(\boldsymbol{z})$ is incurred when the system ends up at state $\boldsymbol{z}$. The recursive value function $v_n(\boldsymbol{r})$ takes the form:

$$v_n(\boldsymbol{r}) = \min_{a \in A(\boldsymbol{r})} \left[ C^a_{\boldsymbol{r}} + \sum_{\boldsymbol{z} \in \Re} P^a_{\boldsymbol{rz}} v_{n-1}(\boldsymbol{z}) \right], \qquad \boldsymbol{r} \in \Re, \qquad n = 1, 2, \ldots$$

If the associated Markov chain is aperiodic (a condition satisfied in our case), then it is possible to obtain upper and lower bounds on the minimal expected costs $g^*$ per transition, and these bounds converge to $g^*$. Thus, starting with $v_0(\boldsymbol{r}) = 0$ for every $\boldsymbol{r} \in \Re$, we compute in each stage the value function $v_n(\boldsymbol{r})$, the lower bound $L_1(n) = \min_{\boldsymbol{r} \in \Re}[v_{n+1}(\boldsymbol{r}) - v_n(\boldsymbol{r})]$ and the upper bound $L_2(n) = \max_{\boldsymbol{r} \in \Re}[v_{n+1}(\boldsymbol{r}) - v_n(\boldsymbol{r})]$. We stop when the difference between $L_2(n)$ and $L_1(n)$ is smaller than a pre-specified tolerance. The value of $[L_1(n) + L_2(n)]/2$ is then an estimation of $g^*$.

In Section 4 we will use this method to compare the results of the heuristics with the optimal solution.Although this technique is widely used to solve large-scale Markov decision problems, from a practical point of view its applicability is limited due to the size of the state and decision spaces. For example, if there were 6 order categories and the maximum demand in each category were 30, we would have at least $30^6 = 729\,000\,000$ possible combinations of values for the vector of requirements $\boldsymbol{r}$ in each period. Furthermore, the exact structure of the optimal strategy can be very complex and therefore unattractive to implement in practice. Therefore, we will focus in the next section on the development of special subclasses of policies with a structure that captures most of the optimal solution's properties, but due to its simplicity, has more intuitive appeal.

# 3 Lot-sizing strategies

Before describing the heuristics for solving our problem, we introduce some of the notation that will be used throughout this section.

Let $\mu_i$ denote the expectation of each random variable $D_i$, that is, $\mu_i = I\!E(D_i) = \sum_{\ell>0} \ell\, d_i(\ell)$ for every $1 \leq i \leq N$, and let $\mu$ be the average demand per period ($\mu = \sum_{i=1}^{N} \mu_i$).

At the end of an arbitrary period $t$, the expected demand for period $t+i$ under the assumption that no required deliveries for period $t+i$ have been produced, is denoted by

$$e_i = \sum_{m=i}^{N} I\!E(D_m) = \sum_{m=i}^{N} \mu_m, \quad i = 1, 2, \ldots, N. \tag{1}$$

Finally, let us define $b_{i\ell}$ as the probability that during the last $i$ periods, a total number of $\ell$ orders was placed for the current period, that is

$$b_{i\ell} = I\!P(D_1 + \ldots + D_i = \ell) = \sum_{\mathcal{L}_{i\ell}} \prod_{k=1}^{i} d_k(\ell_k), \quad i = 1, 2, \ldots, N,\ \ell = 0, 1, \ldots \tag{2}$$

with $\mathcal{L}_{i\ell}$ denoting the set of all possible one-period demands $(\ell_1, \ell_2, \ldots, \ell_N)$ for which the sum of the first $i$ components equals $\ell$.

## 3.1 The $(x, T, w)$-rule

By examining the optimal production policy for some instances of our problem, we could observe that in many cases the decision whether or not to replenish the inventory in a certain period only depends on the number of orders that should be delivered in that period. Moreover, the total amount of finished items covers not only the demand in $r_1$ but also the known orders for some future periods, plus an extra quantity which protects the system against unforeseen short term demand. This means that upon production, two types of finished goods stock are generated: one is a staging post for known orders awaiting delivery (a so-called *dedicated* stock), while the other is a *non-dedicated* stock of items which are not committed to any specific customer orders. Furthermore, we noticed that generally, a production lot covers the same number of known period's demands, and that the amount manufactured for the non-dedicated stock is fixed. To give an example of an optimal policy exhibiting such a structure, consider $N = 4$, $K = 20$, $h = 1$ and $p = 3$. The demand in each order category $i$ follows a binary distribution such that $d_i(0) = 0.5 = d_i(1)$ $(1 \leq i \leq 4)$. The basic outcome is that for vectors of requirements $\boldsymbol{r}$ with $r_1 > 2$, we produce the known demand for the first 4 periods and 3 additional items. The optimal strategy follows this rule in 98% of the states. Only in the states with unexpectedly high or low demand for the periods 2 until 4 we find deviations from the basic rule.

The above structure of the optimal policy results from the interaction between the cost parameters and the demand pattern. For low $r_1$-values postponing production is cheap. If $r_1$ becomes bigger production can no longer be delayed and it is natural to use the available partial information on future order requirements to decide on the amount to be manufactured. By building a dedicated stock, the timely delivery of the orders covered by that stock is ensured and the next setup is delayed. The latter benefit is obtained especially when the holding costs are relatively low since in such a situation it may compensate to cover a larger number of period's demands. Since future demand is not completely known beforehand it also seems logical to maintain a stock of uncommitted items. The production of extra items will help to cope with the arrival of short term orders without having to initiate production again.

Based on these observations, we propose a simple lot-sizing strategy called the $(x, T, w)$-rule. Such rule states that production does not occur in a period with $r_1 < x$. Whenever $r_1 \geq x$, the known demands for the next $T$ periods are manufactured plus $w$ units for the non-dedicated stock. The parameters $x$, $T$ and $w$ are decision variables. If $T = 1$, we obtain the well known $(s, S)$ policy with $s = -x$ and $S = w$. The special case of $w = 0$ corresponds to a decision strategy that applies to make-to-order systems, that is, to systems where production is only carried out to respond directly to known customer orders. Under various conditions, Dellaert [2] and Melo [12] studied in detail such systems.

By making the distinction between dedicated and non-dedicated stock it is possible to make an exact analysis of the $(x, T, w)$-rule. Upon production we immediately update the state vector $\boldsymbol{r}$ by replenishing the inventory with $w$ units and removing the demands that will certainly be satisfied on time. This implies that the action $a = \sum_{i=1}^{T} r_i + w$ generates a new state $\boldsymbol{z} = \boldsymbol{Q}_a(\boldsymbol{r}) + (\ell_1, \ldots, \ell_N)$ with $\boldsymbol{Q}_a(\boldsymbol{r}) = (-w, 0, \ldots, 0, r_{T+1}, \ldots, r_N, 0)$. In addition, the costs of a production decision are given by $K + h\sum_{j=2}^{T}(j-1)r_j + hw$. For large values of $T$ it may be profitable to keep more items as non-dedicated stock, instead of reserving part of the production lot to the orders in $r_2, \ldots, r_T$. This strategy, however, is much more complex to analyze and may also be inconvenient from a practical point of view, since upon production it will not be possible to immediately inform customers that their orders will be met on time, except those with required deliveries in the production period. Therefore, we will keep our 'direct payment' of the holding costs for the reserved items.

Since demand is stationary and the $(x, T, w)$-rule is a fixed production strategy, there is a natural regenerative point in the stochastic process representing $r_1$. As shown by Melo [12], such a regenerative point coincides with a production epoch. Each time production takes place, a cycle is said to begin. In order to determine the expected costs during a cycle we describe the demand requirements in each period by a Markov process with states $(i, j)$ such that

$$\begin{aligned} i &\equiv \text{number of periods elapsed since the last production period,} \\ j &\equiv \text{value of } r_1 \text{ in the } i\text{th period since previous production.} \end{aligned}$$

Observe that production is triggered simply by the value of $r_1$. In addition, the production decision is always the same when $r_1 \geq x$, namely the known demands for the first $T$ periods are manufactured along with $w$ items. On the other hand, it is not difficult to show that the effects of producing during some period will influence the net requirements over the subsequent $T$ periods only. Therefore, we can restrict the state space to $i = 1, \ldots, T-1$ and $j = -w, \ldots, 0, \ldots, x$. In every state $(i, x)$ we aggregate all the cases with $r_1 \geq x$ in the $i$th period. Also, each state $(T-1, j)$ aggregates all the states not only during period $T-1$

but also in subsequent periods. As already mentioned, the expected costs in these states do not involve the time of the last production. In Figure 1 the structure of the Markov chain is displayed.

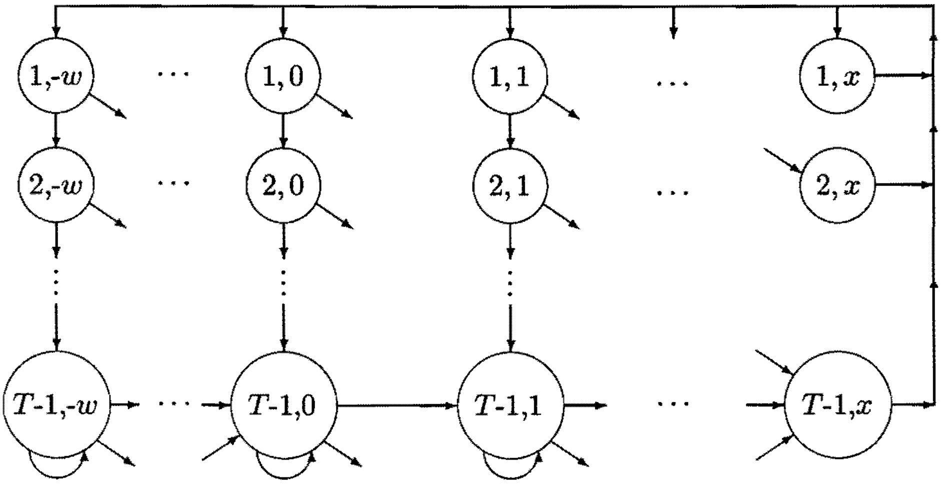


Figure 1: The states and possible transitions in the Markov chain of the $(x, T, w)$-rule for $T \geq 2$.

Let $c_{ij}$ denote the expected costs in each state $(i, j)$ of the above chain. Whenever $j < x$ there is no production. In such case, if $j < 0$ this means that $-j$ items are held in stock with holding costs $-hj$. However, if $j \geq 0$ penalty costs $pj$ are incurred. In the remaining states $(i, x)$, setup costs and usually also holding costs are charged. To avoid a complicated state space, the holding costs corresponding to the known demand are paid immediately. Their expected value depends on the time elapsed since the last production period, that is, on the $i$-value of the state $(i, x)$. For $T > 2$ it holds that if we have just produced in the previous period, the expected number of orders for future periods will not be as large as they would be after $T - 1$ periods, because some of the orders for these periods have already been produced. Since the non-dedicated stock is replenished with $w$ items, holding costs must also be taken into account. In summary, we obtain

$$c_{ij} = \begin{cases} -hj & \text{if } -w \leq j < 0, \\ pj & \text{if } 0 \leq j < x, \\ K + h\sum_{k=1}^{T-1} ke_{k+1} - h\sum_{k=1}^{T-i-1} ke_{k+i+1} + hw & \text{if } j = x. \end{cases}$$

with $e_{k+1}$ and $e_{k+i+1}$ defined by (1). For $j = x$, $c_{ij}$ simplifies into $K + h\sum_{k=1}^{T-1} e_{k+1} \min(i, k)$ $+hw$.

Let $q_{ij}$ denote the proportion of time spent in state $(i, j)$ between two consecutive production periods. Due to the special structure of the Markov chain, the values of $q_{ij}$ can be calculated as follows. If $T \geq 3$, in the first period after a production period we enter state $(1, x)$ or one of the states $(1, j)$. Hence,

$$q_{1j} = \begin{cases} b_{1,j+w} & \text{if } -w \leq j < x, \\ 1 - \sum_{k=-w}^{x-1} b_{1,k+w} & \text{if } j = x. \end{cases}$$

with $b_{1,\cdot}$ given by (2).

For $2 \leq i \leq T-2$, we move to state $(i, j)$ from a state $(i-1, k)$ with probability $b_{i,j-k}$, and enter state $(i, x)$ if the number of orders was less than $x$ in period $i-1$ and at least $x$ in period $i$. This yields

$$q_{ij} = \begin{cases} \sum_{k=-w}^{j} q_{i-1,k}\, b_{i,j-k} & \text{if } -w \leq j < x, \\ \sum_{k=-w}^{x-1} [q_{i-1,k} - q_{ik}] & \text{if } j = x. \end{cases}$$

During one cycle, state $(T-1, j)$ aggregates the visits related to period $T-1$ and all subsequent periods. We can move to this state from state $(T-2, k)$ as well from a state $(T-1, k)$ with $1 \leq k \leq j$. If during the $N$ previous periods no orders have arrived with this specific delivery date, then we stay in the same state with probability $b_{N0}$. Furthermore, there is exactly one production period during one cycle. Hence, if there is no production during the first $T-2$ periods, then it will certainly occur in one of the later periods. Therefore, we have for $T \geq 3$

$$q_{T-1,j} = \begin{cases} \frac{\sum_{k=-w}^{j} q_{T-2,k}\, b_{T-1,j-k} + \sum_{k=-w}^{j-1} q_{T-1,k}\, b_{N,j-k}}{1-b_{N0}} & \text{if } -w \leq j < x, \\ 1 - \sum_{i=1}^{T-2} q_{ix} & \text{if } j = x. \end{cases}$$

For the special case of $T = 2$, the values of $q_{1j}$ are given by

$$q_{1j} = \begin{cases} \frac{b_{1,j+w} + \sum_{k=-w}^{j-1} q_{1k} b_{N,j-k}}{1-b_{N0}} & \text{if } -w \leq j < x \\ 1 & \text{if } j = x \end{cases}$$

while for $T = 1$ we simply need to replace the term $b_{1,j+w}$ in the above expression by $b_{N,j+w}$.

With the knowledge of each $q_{ij}$ the calculation of the expected cycle length becomes straightforward:

$$\mathbb{E}L(x, T, w) = \begin{cases} \sum_{j=-w}^{x} q_{1j} & \text{if } T = 1, \\ \sum_{i=1}^{T-1} \sum_{j=-w}^{x} q_{ij} & \text{if } T \geq 2. \end{cases} \tag{3}$$

In addition, the expected costs during a cycle are determined by

$$\mathbb{E}C(x, T, w) = \begin{cases} \sum_{j=-w}^{x} c_{1j} q_{1j} & \text{if } T = 1, \\ \sum_{i=1}^{T-1} \sum_{j=-w}^{x} c_{ij} q_{ij} & \text{if } T \geq 2. \end{cases} \tag{4}$$

Therefore, the average costs per period $g(x, T, w)$ are obtained by dividing (4) by (3).

To find the triplet that yields the lowest costs per period, we apply certain bounds that allow us to limit the number of triplets that need to be considered. The following result establishes an upper bound on $x^*$ for given values of $T$ and $w$, and is a natural generalization of a bound proposed by Dellaert [2, Chapter 4] for the special case $w = 0$.

**Proposition 1** *For fixed values of $T$ and $w$, the optimal value $x^*$ satisfies*

$$x^* \leq \left\lfloor \frac{g(x^*, T, w)}{p} \right\rfloor + 1.$$

By combining the above bound with the fact that $x^* \geq 1$, we obtain an interval containing the optimal choice for $x$, with $T$ and $w$ fixed.

It is possible to derive better upper bounds for the optimal values of $T^*$ and $w^*$ than the obvious $T^* \leq N$ and $w^* \leq \lfloor \frac{K}{h} \rfloor + 1$ (see also Melo [12]), but determining these bounds is almost as complicated as finding the best $x$, $T$ and $w$ values. Usually, the following procedure offers a quick result:

**Step 0.** Set $x := \lfloor K/p \rfloor + 1$, $T := 1$ and $w := \lfloor K/h \rfloor + 1$.
**Step 1.** Determine $g(x, T, w)$.
`If` $x > \lfloor \frac{g(x,T,w)}{p} \rfloor + 1$ `Then` $x := \lfloor \frac{g(x,T,w)}{p} \rfloor + 1$ and repeat **Step 1**.
**Step 2.** `If` $w > \lfloor \frac{g(x,T,w)}{h} \rfloor + 1$ `Then` $w := \lfloor \frac{g(x,T,w)}{h} \rfloor + 1$.
Determine $K(T) = \min_{x,w} g(x, T, w)$ by first decreasing $w$ as long as $g(x, T, w)$ decreases, and then by decreasing $x$ until no further decrease of $g(x, T, w)$ occurs.
`If` $x$ and $w$ have different values from the ones at the beginning of this step `then` repeat **Step 2**.
**Step 3.** `If` $T = 1$ `or` $(K(T) < K(T-1)$ `and` $T < N)$ `Then` $T := T + 1$
repeat **Step 2**
`Else` Stop.

Observe that for fixed $T$ and decreasing $x$ or $w$, the determination of the average costs does not involve the calculation of every $q_{ij}$ since most of these values are not affected by changes in the components $x$ or $w$. For example, given the triplet $(x, T, w)$ and the corresponding $q_{ij}$-values, we only need to calculate $q_{i,x-1}$ for $1 \leq i \leq T-1$ to be able to determine the costs of the triplet $(x-1, T, w)$. Similarly, we can compute the average costs of the $(x, T, w-1)$-rule by using the information of the triplet $(x, T, w)$. Denoting by $q_{ij}(w)$ the $q_{ij}$-values of the latter rule, it is easy to see that $q_{ij}(w-1) = q_{i,j-1}(w)$ for $1 \leq i \leq T-1$ and $-w+1 \leq j \leq x-1$. Hence, we simply need to calculate $q_{i,x}(w-1)$ for $1 \leq i \leq T-1$. Finally, given the triplet $(x, T-1, w)$ and the corresponding $q_{ij}(T-1)$-values, it is required to determine $q_{T-2,j}(T)$ and $q_{T-1,j}(T)$ for $-w \leq j \leq x$, with respect to the triplet $(x, T, w)$ and in case $T \geq 4$. Otherwise, the new $q_{ij}(T)$ have to be recalculated completely.

## 3.2 The least-cost-per-period strategy

For a single product inventory problem subject to a deterministic time-varying demand pattern, Silver & Meal [14] designed a lot-sizing heuristic which consists of selecting the replenishment amount that produces the (first local) minimum of the total relevant costs per unit of time. These costs are obtained by dividing the setup cost and the total holding costs of a replenishment, by the number of periods corresponding to the demand requirements that are covered by the replenishment. This strategy is myopic in that it looks at the replenishment in the current period only and not at its relation to future replenishments. The rule is simple, intuitively appealing and seems to perform well in many situations as reported by Blackburn & Millen [1] and Silver, Pyke & Peterson [15, Chapter 6]).

In this section we derive a least-cost-per-period (LCP) strategy by applying the same criterion as in the Silver-Meal heuristic and taking into account that backlogs are permitted in our case.

As in the $(x, T, w)$-rule, we consider that whenever production occurs, part or whole of the known demand requirements is manufactured together with an extra amount for a buffer stock. However, in contrast to the previous rule, the new strategy will attempt to adapt to the pattern of known customer orders. This means that we consider each state vector $\boldsymbol{r}$ separately and restrict our attention to a subclass of production decisions of the action space $A(\boldsymbol{r})$ that are characterized by two components, $T$ and $w$. The first component indicates the number of periods with known demands that are covered by a batch, while the second component gives the extra number of items that are included in a production lot to replenish the non-dedicated stock. Using this approach, the direct costs incurred by selecting action $a = \sum_{i=1}^{T} r_i + w$ on observing state $\boldsymbol{r}$ are determined by

$$C_{\boldsymbol{r}}^{a} = \begin{cases} -hr_1 & \text{if } a = 0 \text{ and } r_1 \leq 0, \\ pr_1 & \text{if } a = 0 \text{ and } r_1 > 0, \\ K + h\sum_{i=1}^{T-1} ir_{i+1} + hw & \text{if } a \geq r_1. \end{cases}$$

Apart from the above costs we will also consider indirect costs in the decision criterion by assuming that if the known requirements for the next $T$ periods are covered, then the following production will be planned $T$ periods from now. This assumption is necessary since there is no production trigger as in the $(x, T, w)$-rule. Its impact is measured by extra penalty costs during $T - 1$ periods without production. Due to the possibility of building a stock of uncommitted items, we may also have to pay holding costs. Consequently, the costs charged during the $T - 1$ periods immediately after a production in some period $t$, depend on the initial amount $w$ in the non-dedicated stock and also on the number of orders that will be placed during periods $t, t+1, \ldots, t+T-2$ with due dates in $t+1, \ldots, t+T-1$. Since this future demand is unknown at $t$, we use its expected value. To give an example, suppose that $N = 4$ and we produce during $t$ the known orders until $t + 2$ (i.e. $T = 3$) plus $w$ items. The demand that arrives during $t$ to be delivered by the end of $t + 1$ has an expectation $\mu_1$. If the amount $w$ in stock is larger than $\mu_1$, we will pay holding costs $h(w - \mu_1)$, otherwise we will have penalty costs $p(\mu_1 - w)$. In each case, at the end of period $t+1$ the inventory level becomes $w - \mu_1$. The expected demand for period $t+2$ that arrives during $t$ and $t+1$, is given by $\mu_2$ and $\mu_1$, respectively. As a result, if $\mu_1 + \mu_2 \leq w - \mu_1$, the expected demand for period $t + 2$ can be satisfied and the remaining items in stock must pay holding costs. Otherwise, penalty costs are charged.

Summarizing, for the action $a = \sum_{i=1}^{T} r_i + w$ with $T \geq 2$, the holding and penalty costs during the next $T-1$ periods given an initial buffer stock of $w$ items, are estimated by

$$\begin{aligned} HP(T,w) &= h\sum_{i=2}^{T}\left[w-\sum_{k=2}^{i-1}\sum_{j=1}^{k-1}\mu_j-\sum_{j=1}^{i-1}\mu_j\right]^+ + \\ & \quad p\sum_{i=2}^{T}\left[-w+\sum_{k=2}^{i-1}\sum_{j=1}^{k-1}\mu_j+\sum_{j=1}^{i-1}\mu_j\right]^+ . \end{aligned}$$

with $[m]^+ = \max(0,m)$. Clearly, $HP(T,w) = 0$ for $T = 0,1$.

For the domain of $w$, we choose values depending on $T$ in such a way that the amount that can be manufactured for the non-dedicated stock is bounded from above by the total expected demand for the periods without production. In short, $w \in \{0,\ldots,\sum_{i=2}^{T}\sum_{j=1}^{i-1}\mu_j\}$.

Hence, on observing a state $\boldsymbol{r}$ we select the action $a' = \sum_{i=1}^{T} r_i + w$ that minimizes the expected costs per period, that is

$$a' = \arg\min_a \left\{\frac{C_{\boldsymbol{r}}^a + HP(T,w)}{\max\{1,T\}}\right\}$$

with $T \in \{0,1,\ldots,N\}$ and $w \in \{0,\ldots,\sum_{i=2}^{T}\sum_{j=1}^{i-1}\mu_j\}$.

The long-run average costs per period of the LCP approach are obtained by applying a similar iteration scheme to that of the optimal lot-sizing policy or by simulation, if the number of possible states is too large. Observe that although the choice of the action $a'$ is influenced by the costs incurred during those periods without production, in practice we may produce again, even in the following period.

## 3.3 The $(R,S)$ policy

In a classical $(R,S)$ system, every $R$ units of time the inventory level is monitored and a sufficient amount is ordered to raise the inventory position to the level $S$ (Hax & Candea [10, Chapter 4]). The $(R,S)$ policy (or periodic replenishment rule) is a simple strategy that is largely used in practice. Applying this rule to our problem implies that a setup occurs every $R$ periods, unless there is no new demand for the next $R$ periods and the stock level is still $S$. To raise the inventory level to $S$ at each review epoch, a total of $r_1+S$ items must be manufactured. As a result, the costs $K+hS$ are charged at the end of the production period. However, since production will not take place in the $R-1$ periods following a setup, we also need to estimate the holding and penalty costs during those periods. Let $L(y)$ denote the one-period expected holding and penalty costs when the stock on hand minus backorders is $y$, and the demand for the period is not included. The function $L(y)$ often appears in the analysis of classical inventory problems (Veinott & Wagner [17]) and is defined by

$$\begin{aligned} L(y) &= h\sum_{k=0}^{y}(y-k)b_{Nk} + p\sum_{k=y+1}^{\infty}(k-y)b_{Nk}, \qquad y = \ldots,-1,0,1,\ldots \\ &= (h+p)\sum_{k=0}^{y}(y-k)b_{Nk} + p(\mu - y), \qquad y = \ldots,-1,0,1,\ldots \qquad (5) \end{aligned}$$

In case $R \geq 2$ and production takes place in some period $t+1$, then the expected costs in period $t+2$ are given by $L(S)$. In each one of the following periods $t+i$ with $i=3,\ldots,R$, the costs partly depend on the orders for periods $t+2,\ldots,t+i-1$. The total demand for one period is given by the convolution of the $N$ variables $D_1,\ldots,D_N$. Hence, the total demand for say $i$ periods is obtained by convoluting $iN$ random variables. As a result, the expected costs incurred in those $R-1$ periods without production are given by

$$G(R,S) = L(S) + \sum_{i=1}^{R-2}\sum_{k=0}^{\infty} b_{Nk}^{(i)} L(S-k), \quad R \geq 2, \tag{6}$$

with $b_{Nk}^{(i)}$ the $i$-fold convolution of $b_{Nk}$.

Observe that if there are no demand requirements in the $R$ periods immediately following an inventory replenishment, then at the beginning of the next cycle the inventory position is still at $S$. Consequently, production does not occur and a setup cost is not incurred. Since the probability that there is demand for at least one of the next $R$ periods is given by $(1-(b_{N0})^R)$, it follows by (6) that for fixed $R$ and $S$, the average costs per period are determined by

$$g(R,S) = \frac{K(1-(b_{N0})^R) + hS + G(R,S)}{R}.$$

In order to find the pair $(R,S)$ with the lowest average costs, we make use of the fact that for $R$ fixed, the function $F(S) = hS + G(R,S)$ is convex (since $L(y)$ is a convex function). Therefore, for fixed $R \geq 2$, the best choice for $S$ is the unique integer satisfying

$$\Delta F(S-1) < 0 \leq \Delta F(S),$$

where $\Delta F(y) = F(y+1) - F(y)$. Due to the form of the cost function $L(y)$ in (5), it can easily be seen that $L(y) = L(y-1) + (h+p)\sum_{k=0}^{y-1} b_{Nk} - p$ and so by (6) it follows that

$$\begin{aligned}
G(R,S+1) - G(R,S) &= (h+p)\sum_{k=0}^{S} b_{Nk} - p - (R-2)p + \\
&\quad (h+p)\sum_{i=1}^{R-2}\sum_{k=0}^{\infty}\sum_{j=0}^{S-k} b_{Nk}^{(i)} b_{Nj} \\
&= (h+p)\sum_{k=0}^{S} b_{Nk} - (R-1)p + (h+p)\sum_{i=1}^{R-2}\sum_{k=0}^{S} b_{Nk}^{(i+1)} \\
&= (h+p)\sum_{i=1}^{R-1}\sum_{k=0}^{S} b_{Nk}^{(i)} - (R-1)p.
\end{aligned}$$

Therefore, $S$ is determined from

$$\sum_{i=1}^{R-1}\sum_{k=0}^{S-1} b_{Nk}^{(i)} < \frac{(R-1)p-h}{p+h} \leq \sum_{i=1}^{R-1}\sum_{k=0}^{S} b_{Nk}^{(i)}. \tag{7}$$

It is clear that for $R = 1$ the best value for $S$ is 0 and the corresponding average costs are given by $K(1 - b_{N0})$. Starting with $R = 1$, we set $R = R+1$, find the $S$-value satisfying the inequalities in (7), and calculate the average costs of the new pair. After each increment of $R$, the lowest costs obtained until then are saved. This procedure is repeated until we obtain a pair with higher costs than the best pair so far. In all the numerical tests reported in the next section, this stopping criterion yielded the optimal pair $(R^*, S^*)$.

# 4 Numerical results

In this section we present the results of our computational experiences. First we consider a small test set, where we compare the results of the heuristics with the optimal solution obtained by the Markov formulation of Section 2. Then we perform a large set of tests covering variations in both the cost and the demand parameters.

## 4.1 Comparison with optimal solution

The optimal solution can only be calculated for small problems. For larger problems the state space grows too large. Therefore we only consider 16 cases with a binary demand and a horizon of 5 periods. As a measure of the variability of the demand, we consider the coefficient of variation $cv$ which is defined as the ratio of the standard deviation $\sigma$ to the mean total demand per period $\mu$. Observe that since the holding cost is normalized to one, the setup and penalty costs in Table 1 are in fact ratios of those actual costs to the holding cost

From Table 1 we can observe that on average, the difference between the costs of the $(x, T, w)$-rule and the optimal costs is less than 1 %. The corresponding values for the LCP strategy, the $(s, S)$ policy and the $(R, S)$ policy are 2.6 %, 3.9 % and 16.5 % respectively.

## 4.2 Comparison of heuristics

In this subsection we consider a large test set as shown in Table 2. In total, 324 cases were examined covering a sufficiently wide variety of combinations to offer insight into the quality of the heuristic procedures presented in the previous section.

For values of $\mu$ and $cv$ such that $\sigma^2/\mu < 1$, the demand $D_i$ in each order category is assumed to follow a binomial distribution. In those cases with $\sigma^2/\mu \geq 1$, $D_i$ is described by a negative binomial distribution. For $\mu$ and $N$ fixed, the mean $\mu_i$ of the variable $D_i$ $(1 \leq i \leq N)$ is obtained by first generating numbers $v_i$ from a uniform distribution in $[0, 1)$, and then calculating $\mu_i = \mu v_i / \sum_{j=1}^{N} v_j$. With the values of $\mu_i$, the corresponding parameters in the binomial and negative binomial distributions are determined.

The performance of the $(x, T, w)$-rule is examined not only through the best triplet $(x^*, T^*, w^*)$, but also by simply approximating the decision variables by means of a rule-of-thumb. The well-known EOQ model for constant demand (Silver, Pyke & Peterson [15, Chapter 5]) provides us with a simple way for deriving approximations for the values of $x$

| cv | K | p | $(x^*, T^*, w^*)$ | gap | $(s^*, S^*)$ | gap | LCP-gap | $(R^*, S^*)$ | gap |
|---|---|---|---|---|---|---|---|---|---|
| 0.89 | 8 | 3 | (2,3,1) | 1.97 | (-2,2) | 6.33 | 2.72 | (4,2) | 32.88 |
| 0.89 | 8 | 6 | (1,3,1) | 0.34 | (-1,3) | 5.31 | 3.26 | (4,3) | 46.23 |
| 0.89 | 32 | 3 | (3,5,4) | 0.21 | (-3,6) | 1.46 | 11.90 | (7,5) | 19.47 |
| 0.89 | 32 | 6 | (2,5,4) | 0.00 | (-2,6) | 1.31 | 6.53 | (7,6) | 24.41 |
| 0.55 | 8 | 3 | (2,2,1) | 1.85 | (-2,3) | 4.34 | 4.49 | (3,3) | 18.55 |
| 0.55 | 8 | 6 | (1,2,1) | 2.46 | (-1,4) | 5.00 | 3.32 | (3,4) | 25.83 |
| 0.55 | 32 | 3 | (4,4,5) | 0.28 | (-4,8) | 1.18 | 1.72 | (6,8) | 9.20 |
| 0.55 | 32 | 6 | (2,5,5) | 0.11 | (-2,9) | 1.03 | 0.96 | (6,9) | 14.05 |
| 0.37 | 8 | 3 | (2,2,1) | 0.88 | (-2,4) | 5.46 | 0.92 | (3,5) | 13.75 |
| 0.37 | 8 | 6 | (1,2,1) | 0.83 | (-2,4) | 8.71 | 0.75 | (2,4) | 18.10 |
| 0.37 | 32 | 3 | (4,4,5) | 0.44 | (-4,10) | 1.12 | 1.26 | (5,10) | 5.24 |
| 0.37 | 32 | 6 | (2,5,5) | 0.38 | (-3,11) | 1.01 | 0.10 | (5,11) | 7.80 |
| 0.22 | 8 | 3 | (3,2,1) | 0.40 | (-3,4) | 6.89 | 0.79 | (2,4) | 8.40 |
| 0.22 | 8 | 6 | (2,2,1) | 0.32 | (-2,5) | 11.00 | 0.55 | (2,5) | 13.38 |
| 0.22 | 32 | 3 | (5,4,5) | 0.10 | (-5,12) | 0.72 | 1.47 | (5,13) | 2.27 |
| 0.22 | 32 | 6 | (3,4,5) | 0.12 | (-3,13) | 1.55 | 0.10 | (4,12) | 4.17 |

Table 1: Percentual extra costs compared to the optimal policy.

and $T$. For the number of periods whose known demands are covered by a production lot, we select

$$\tilde{T} = \min\left\{N, \left\lfloor \sqrt{\frac{2K}{h\mu}} + 0.5 \right\rfloor\right\}.$$

For the above value of $\tilde{T}$, it follows that the average costs per period according to the EOQ model, are given by $C(\tilde{T}) = K/\tilde{T} + \mu\tilde{T}h/2$. Hence, by applying the upper bound on the value of $x$ provided by Proposition 1, we obtain the following approximation:

$$\tilde{x} = \left\lfloor \frac{C(\tilde{T})}{p} \right\rfloor + 1. \tag{8}$$

Our choice of $\tilde{T}$ influences the amount that must be produced to replenish the non-dedicated stock. A logical approximation consists in taking the expected demand that can be placed after the current period for delivery in the next $\tilde{T}$ periods, and that was not included in the production lot, i.e. $\sum_{i=2}^{\tilde{T}} \sum_{j=1}^{i-1} \mu_j$. However, if both the setup and penalty costs are relatively large and the unit holding cost is low, it may be advantageous to manufacture more than the above quantity. Observe that by building a dedicated stock with the known demands in the first $\tilde{T}$ periods, we do not cover any requirements beyond $\tilde{T}$. It is reasonable to assume that in each period after $\tilde{T}$, $\mu$ items need to be delivered. Hence, we compare the EOQ estimation of the total costs per period $C(\tilde{T})$, with the holding costs

| Parameter | Symbol | Values tested |
|---|---|---|
| holding cost per period | $h$ | 1 |
| setup cost per period | $K$ | $100, 200, 500$ |
| penalty cost per period | $p$ | $5, 10, 15$ |
| number of demand categories | $N$ | $4, 5, 6$ |
| mean demand per period | $\mu$ | $5, 10, 50$ |
| coefficient of variation | $cv$ | $0.25, 0.5, 1, 1.5$ |

Table 2: Combination of parameters tested.

that should be paid if we would keep $\mu$ units in stock in every period. This yields the following estimation for the value of $w$:

$$\widetilde{w} = \sum_{i=2}^{\widetilde{T}} \sum_{j=1}^{i-1} \mu_j + (n - \widetilde{T})\mu \tag{9}$$

with $n = \max\left(\widetilde{T}, \lfloor \frac{C(\widetilde{T})}{\mu h} \rfloor\right)$.

Since the $(s, S)$ policy is a special case of the $(x, T, w)$-rule, we apply the approximation in (8) to obtain $\widetilde{s} = -\widetilde{x}$. In addition, we estimate the size of a production lot by the corresponding value in the EOQ model, that is

$$\widetilde{S} - \widetilde{s} = \sqrt{\frac{2K\mu}{h}}.$$

Table 3 presents the average costs per period obtained by each heuristic procedure over the 324 test problems. These costs result from adding the average setup, holding and penalty costs indicated in columns 3, 4 and 5. An important performance criterion in a production system is the service level, that is, the percentage of customer orders delivered on time. This value is given in column 6.

From Table 3 we observe that both in the $(x^*, T^*, w^*)$-rule and the $(s^*, S^*)$ policy the number of occasions with production is approximately the same. The $(s^*, S^*)$ policy cannot adapt to the demand requirements as well as the $(x^*, T^*)$-rule, and therefore incurs higher holding and penalty costs. The LCP strategy seems to follow a different approach by initiating production frequently and by replenishing inventory with less items than in the $(x^*, T^*, w^*)$ and $(s^*, S^*)$ policies. Furthermore, by inspecting the penalty costs, it appears that with the LCP strategy the system is less protected against stockouts. This latter aspect is more evident in the $(R^*, S^*)$ policy, as expected.

To study the performance of the lot-sizing strategies in detail, we divided the average costs of each rule by the costs of the optimal triplets $(x^*, T^*, w^*)$. Table 4 indicates the distribution of the percentage of deviations. The 'Average % Dev.' column presents the average of the deviations. The next block of eight columns gives the distribution of

| Strategy | average costs | setup | holding | penalty | serv(%) |
|---|---|---|---|---|---|
| $(x^*,T^*,w^*)$-rule | 80.72 | 43.83 | 32.78 | 4.11 | 96.73 |
| $(\tilde{x},\tilde{T},\tilde{w})$-rule | 83.19 | 40.37 | 36.70 | 6.12 | 95.60 |
| $(s^*,S^*)$ policy | 82.73 | 44.30 | 34.10 | 4.33 | 96.60 |
| $(\tilde{s},\tilde{S})$ policy | 85.62 | 37.93 | 41.46 | 6.23 | 94.80 |
| LCP | 86.64 | 52.48 | 26.87 | 7.29 | 92.98 |
| $(R^*,S^*)$ policy | 107.53 | 58.20 | 35.63 | 13.70 | 85.43 |

Table 3: Average performance of the strategies (324 cases).

the percentage of deviations for each heuristic. For example, in 20% of the cases, the LCP strategy was at least 0.5% better than the $(x^*,T^*,w^*)$-rule.

| Strategy | Average % Dev. | Percentage of Cases Whose Deviation Was Strictly Less than | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | −0.5% | 0% | 0.5% | 1% | 5% | 10% | 20% | 40% |
| $(\tilde{x},\tilde{T},\tilde{w})$-rule | 3.3 | 0 | 0 | 27 | 42 | 81 | 91 | 96 | 100 |
| $(s^*,S^*)$ policy | 2.3 | 0 | 0 | 34 | 45 | 87 | 96 | 100 | 100 |
| $(\tilde{s},\tilde{S})$ policy | 5.8 | 0 | 0 | 10 | 24 | 66 | 81 | 94 | 100 |
| LCP | 9.1 | 20 | 30 | 40 | 46 | 60 | 71 | 82 | 94 |
| $(R^*,S^*)$ policy | 32.1 | 0 | 0 | 0 | 0 | 18 | 29 | 46 | 67 |

Table 4: Performance of the strategies relative to the $(x^*,T^*,w^*)$-rule (324 cases).

It can be seen that the average costs of the optimal pairs $(s^*,S^*)$ deviate very little from those of the $(x^*,T^*,w^*)$-rule. Also, the approximations of the parameters $x$, $T$ and $w$ yield good results and the approximations $\tilde{s}$ and $\tilde{S}$ do not reveal very large deviations. The LCP strategy performs better than the $(x^*,T^*,w^*)$-rule in 30% of the problems, but there appears to exist cases for which large deviations occur. Finally, the $(R^*,S^*)$ policy performs poorly.

To identify the conditions that most influence the behavior of each rule, we analyzed the results of several parameter combinations and realized that the magnitude of the setup cost and the demand variability are responsible to a large extent to the results displayed in Table 4.

We will start by studying the effect of $K$. In Tables 5 and 6 we compare the distributions of the percentage of deviations for the lowest and the highest values of the setup cost in our test bed. Striking differences in the performance of almost every rule are noticeable when $K$ increases from 100 to 500. With a relatively small setup cost, the LCP strategy is in 38% of the cases better than the $(x^*,T^*,w^*)$-rule. In the remaining problems, the deviations are always less than 12%. The good performance of this strategy is due to the

| Strategy | Average % Dev. | Percentage of Cases Whose Deviation Was Strictly Less than | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | −0.5% | 0% | 0.5% | 1% | 5% | 10% | 20% | 40% |
| $(\tilde{x}, \tilde{T}, \tilde{w})$-rule | 1.1 | 0 | 0 | 38 | 54 | 100 | 100 | 100 | 100 |
| $(s^*, S^*)$ policy | 3.5 | 0 | 0 | 23 | 36 | 76 | 87 | 100 | 100 |
| $(\tilde{s}, \tilde{S})$ policy | 9.0 | 0 | 0 | 9 | 23 | 48 | 64 | 87 | 100 |
| LCP | 1.4 | 20 | 38 | 53 | 64 | 86 | 97 | 100 | 100 |
| $(R^*, S^*)$ policy | 33.3 | 0 | 0 | 0 | 0 | 13 | 19 | 44 | 69 |

Table 5: Performance of the strategies relative to the $(x^*, T^*, w^*)$-rule for small setup cost ($K = 100$; 108 cases).

fact that contrary to the $(x^*, T^*, w^*)$-rule, a decision to produce in a certain period is based not only on the value of $r_1$ but also on the demand requirements in the other components of the order state vector. Moreover, a batch does not always cover the known demands for the same number of periods and the amount manufactured for the non-dedicated stock is not fixed. Thus, the strategy is more flexible than the $(x^*, T^*, w^*)$-rule. However, when $K$ becomes large the quality of the results deteriorates. In that case, it is natural that to prevent from paying too many setup costs, one should restrict the number of occurrences of production, and in every period with production, larger batches should be manufactured. This is particularly important with respect to the amount that replenishes the non-dedicated stock. Recall from Section 3.2 that in the LCP approach, the maximum number of items that can be produced for the non-dedicated stock is bounded from above by the expected demand for those periods covered by the production of known orders. Obviously, such a bound does not depend on the magnitude of the setup cost, but in situations with large $K$ it is often beneficial to increase the production of uncommitted items beyond this limit. Since this is not permitted, the rule tends to initiate production very often at the expense of high setup costs. These features account for the large deviations displayed by the LCP strategy in Table 6 and for a total average deviation of 20.4%. Nevertheless, in 28% of the problems, the strategy still performs better than the $(x^*, T^*, w^*)$-rule. These cases refer mostly to problems with relatively high demand variability. 00000 00000 00000 00000 00000 00000

The $(x^*, T^*, w^*)$-rule tends to overestimate the values of $T$ and $w$ and as a result, the holding costs increase. Although in the LCP strategy production still takes place more often, the extra setups together with lower holding costs yield in total somewhat smaller costs than in the $(x^*, T^*, w^*)$-rule. The approximations $\tilde{x}$, $\tilde{T}$ and $\tilde{w}$ also seem to be affected by the size of $K$. The values of $\tilde{x}$ overestimate $x^*$ with the resulting effect that higher penalty costs are incurred. With respect to $\tilde{T}$, the estimated values often equal $T^*$ or $T^* + 1$. From the expression in (9), we notice that if $K$ increases then $\tilde{w}$ grows as well. This often results in a very large value which contributes to the deviations provided in Table 6.

| Strategy | Average % Dev. | Percentage of Cases Whose Deviation Was Strictly Less than | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | −0.5% | 0% | 0.5% | 1% | 5% | 10% | 20% | 40% |
| $(\tilde{x}, \tilde{T}, \tilde{w})$-rule | 7.0 | 0 | 0 | 6 | 17 | 56 | 75 | 89 | 100 |
| $(s^*, S^*)$ policy | 1.2 | 0 | 0 | 48 | 63 | 97 | 100 | 100 | 100 |
| $(\tilde{s}, \tilde{S})$ policy | 3.6 | 0 | 0 | 8 | 24 | 81 | 90 | 100 | 100 |
| LCP | 20.4 | 19 | 28 | 31 | 32 | 33 | 39 | 53 | 82 |
| $(R^*, S^*)$ policy | 29.2 | 0 | 0 | 0 | 0 | 24 | 39 | 50 | 69 |

Table 6: Performance of the strategies relative to the $(x^*, T^*, w^*)$-rule for large setup cost ($K = 500$; 108 cases).

| Strategy | Average % Dev. | Percentage of Cases Whose Deviation Was Strictly Less than | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | −0.5% | 0% | 0.5% | 1% | 5% | 10% | 20% | 40% |
| $(\tilde{x}, \tilde{T}, \tilde{w})$-rule | 2.7 | 0 | 0 | 48 | 63 | 84 | 93 | 96 | 100 |
| $(s^*, S^*)$ policy | 2.0 | 0 | 0 | 70 | 78 | 86 | 89 | 100 | 100 |
| $(\tilde{s}, \tilde{S})$ policy | 3.8 | 0 | 0 | 33 | 57 | 78 | 88 | 94 | 100 |
| LCP | 11.3 | 19 | 30 | 41 | 44 | 59 | 65 | 79 | 93 |
| $(R^*, S^*)$ policy | 5.5 | 0 | 0 | 0 | 0 | 68 | 85 | 99 | 100 |

Table 7: Performance of the strategies relative to the $(x^*, T^*, w^*)$-rule for low demand variability ($cv = 0.25$; 81 cases).

In the $(s^*, S^*)$ policy a large setup cost has the reverse effect. This is due to the fact that since only the known requirements for the first period are manufactured, it is necessary to keep a large inventory of items to cope with demands in the following periods. Clearly, the $(s^*, S^*)$ policy builds a larger non-dedicated stock than the $(x^*, T^*, w^*)$-rule. If $K$ is relatively high, however, the differences tend to decrease. This is also observed for the approximations of the reorder point $\tilde{s}$ and the order-up-to level $\tilde{S}$. Finally, the $(R^*, S^*)$ policy shows considerable deviations to the $(x^*, T^*, w^*)$-rule. The demand variability has a stronger impact on the performance of this policy than the magnitude of the setup cost. This is confirmed in Tables 7 and 8 where the distribution of the percentage of deviations for $cv = 0.25$ and $cv = 1.5$ are presented. It is interesting to see that the LCP strategy is sensitive to demand fluctuations and yields better results for large values of $cv$. The performance of all the other strategies deteriorates as the variability of the demand increases.

| Strategy | Average % Dev. | Percentage of Cases Whose Deviation Was Strictly Less than | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | −0.5% | 0% | 0.5% | 1% | 5% | 10% | 20% | 40% |
| $(\tilde{x}, \tilde{T}, \tilde{w})$-rule | 4.7 | 0 | 0 | 0 | 11 | 70 | 88 | 96 | 100 |
| $(s^*, S^*)$ policy | 3.4 | 0 | 0 | 7 | 11 | 75 | 100 | 100 | 100 |
| $(\tilde{s}, \tilde{S})$ policy | 11.1 | 0 | 0 | 0 | 0 | 31 | 60 | 80 | 100 |
| LCP | 6.5 | 26 | 31 | 37 | 37 | 60 | 78 | 88 | 96 |
| $(R^*, S^*)$ policy | 68.4 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 15 |

Table 8: Comparison of the performance of the strategies to the $(x^*, T^*, w^*)$-rule for large demand variability ($cv = 1.5$; 81 cases).

# 5 Conclusions

In this paper we proposed two different strategies to compute near-optimal decisions for a stochastic single-item make-to-stock problem with partial customer-order information. We noticed that for the majority of the possible states the optimal solutions to our lotsizing problem have the same structure as in the so-called $(x, T, w)$-rule, where production is initiated during a period for which the required deliveries are at least $x$ units. In that case, a batch covers the known demand requirements for the next $T$ periods plus $w$ units. The finished items corresponding to the known demands are reserved for delivery at the promised due dates. The remaining items are not committed to any specific customer orders and form a non-dedicated stock. In the least-cost-per-period (LCP) strategy, the production decisions are also based on all available information concerning future customer requirements. However, in contrast to the $(x, T, w)$-rule, lot sizes do not cover the demand for a fixed number of periods. Also, the non-dedicated stock may be replenished with a variable amount. The performance of these two lot-sizing strategies was compared to that of the well-known $(s, S)$ and $(R, S)$ inventory policies. As remarked in Section 3.1, the $(s, S)$ policy is in fact a special case of the $(x, T, w)$-rule with $s = -x$, $T = 1$ and $S = w$. By only covering the requirements in the current period, the $(s, S)$ policy does not explore the structure of the demand process completely. This feature also applies to the $(R, S)$ policy but has a stronger impact on the performance of this rule since it is only allowed to produce at fixed intervals. In that respect, the $(s, S)$ policy is more flexible and can choose the reorder point and the order-up-to level so that the system is not exposed to stockouts very often. This characteristic was confirmed by the extensive numerical tests conducted in Section 4. Nevertheless, taking into account the known required deliveries in the next periods leads to a better planning of production as observed in the $(x, T, w)$-rule, especially when demand has large fluctuations. For problems with both small setup cost and demand variability, the LCP strategy yielded very good results. In the remaining cases, the $(x, T, w)$-rule gave the lowest average costs. We also derived approximations for the decision variables $x$, $T$ and $w$ by applying the EOQ model. Surprisingly, the costs obtained were usually only a few percent higher than those of the optimal values $x^*$, $T^*$ and $w^*$. The approximations of $s^*$ and $S^*$ seemed to be less robust, especially when demand showed large fluctuations.

All presented heuristics could also be used for the alternative problem, described by Heath and Jackson [11] and others, where the demand for the first period is *not* known with certainty. Only slight modifications to the direct costs in the various states would be necessary and undoubtedly, the optimal $(x^*, T^*, w^*)$-solution for such a problem will contain most of the structure of the optimal solution.

# References


[1] J. D. Blackburn and R. A. Millen. Heuristic lot-sizing performance in a rolling-schedule environment. *Decision Sciences*, 11:691–701, 1980.

[2] N. P. Dellaert. Production to order: models and rules for production planning. *Lecture Notes in Economics and Mathematical Systems, no. 333.* Springer-Verlag, Berlin, 1989.

[3] N. P. Dellaert and M. T. Melo. Heuristic procedures for a stochastic lot-sizing problem in make-to-order manufacturing. *Annals of Operations Research*, 59:227–258, 1995.

[4] N. P. Dellaert and M. T. Melo. Production strategies for a stochastic lot-sizing problem with constant capacity. *European Journal of Operational Research*, 92:281–301, 1996.

[5] N. P. Dellaert and J. Wessels. Production scheduling with uncertain demand and no stockkeeping. In *Operations Research Proceedings*, pages 247–252, 1985.

[6] S. C. Graves, D. B. Kletter, and W. B. Hetzel. A dynamic model for requirements planning with application to supply chain optimization. *Operations Research*, 46,3:35–49, 1998.

[7] G. Gregory, S. A. Klesniks, and J. A. Piper. Batch production decisions and the small firm. *Journal of the Operational Research Society*, 34:469–477, 1983.

[8] R. Güllü. On the value of information in dynamic production/inventory problems under forecast evolution. *Naval Research Logistics*, 43:289–303, 1996.

[9] R. Hariharan and P. Zipkin. Customer-order information, leadtimes and inventories. *Management Science*, 41:1599–1607, 1995.

[10] A. C. Hax and D. Candea. *Production and Inventory Management.* Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1984.

[11] D. C. Heath, and P. L. Jackson. Modeling the evolution of demand forecasts with applicaton to safety stock analysis in production/distribution systems. *IIE Trans.*, 26,3:17–30, 1994.

[12] M. T. Melo. *Stochastic Lot-Sizing in Production Planning: Strategies for make-to-order and make-to-stock*, volume 128. Tinbergen Institute Research Series, Thesis Publishers, Amsterdam, The Netherlands, 1996.

[13] A. Odoni. On finding the maximal gain for Markov Decision Processes. *Operations Research*, 17:857–860, 1969.

[14] E. A. Silver and H. C. Meal. A heuristic for selecting lot size requirements for the case of a deterministic time-varying demand rate and discrete opportunities for replenishment. *Production and Inventory Management*, 14:64–74, 1973.

[15] E. A. Silver, D. F. Pyke, and R. Peterson. *Inventory Management and Production Planning and Scheduling.* John Wiley, New York, 1998.

[16] H. C. Tijms. *Stochastic modelling and analysis–An algorithmic approach.* John Wiley, New York, 1994.

[17] A. F. Veinott and H. M. Wagner. Computing optimal (s,S) inventory policies. *Management Science*, 11:525–552, 1965.

[18] J. Wijngaard. *Foreknowledge of demand in combination with a high utilization rate: The single stage, single product case.* Paper presented at the Ninth International Working Seminar on Production Economics, Igls, Austria, February 19-23, 1996.

[19] W. H. M. Zijm. Approximate stochastic dynamic programming: an application to production control. In A. Volgenant, editor, *Caleidoscoop van de Levende Besliskunde*, pages 163–173. Liber Amicorum Prof.Dr. G. de Leve, Amsterdam, 1991 (in Dutch).

BETA

Working Papers

| | | |
|---|---|---|
| WP 1 | The Dutch IOR approach to organisational design : an alternative to business process reengineering<br>**Frans M. van Eijnatten and Ad H. van der Zwaan** | |
| WP 2 | Simultaneous layout planning, organisation design and technology development : preparing for world class standards in production<br>**Ronald J.H. van de Kuil, Frans M. van Eijnatten, Andre G.W. Peeters and Maarten J. Verkerk** | |
| WP 3 | A framework for implementation of statistical process control<br>**R.J.M.M. Does, A. Trip, W.A.J. Schippers** | |
| WP 4 | Bounds for performance characteristics : a systematic approach via cost structures<br>**G.J. van Houtum, W.H.M. Zijm, I.J.B.F. Adan, J. Wessels** | |
| WP 5 | Vertekeningen in budgetten : oorzaak, gevolg en probleemaanpak<br>**Jeroen Weimer** | |
| WP 6 | Corporate governance en financiele ondernemingsdoelen : winst, profit en Gewinn<br>**Jeroen Weimer en Joost C. Pape** | |
| WP 7 | A behavioral view on corporate governance and corporate financial goals of Dutch, US, and German firms<br>**Jeroen Weimer and Joost C. Pape** | |
| WP 8 | Corporate governance : het belang van theorie voor de praktijk<br>**Joost C. Pape en Jeroen Weimer** | |
| WP 9 | Mixed policies for recovery and disposal of multiple type assembly products : commercial exploitation of compulsory return flows<br>**H.R. Krikke, P.C. Schuur and A. van Harten** | |
| WP 10 | Balancing stocks and flexible recipe costs and high service level requirements in a batch process industry : a study of a small scale model<br>**W.G.M. Rutten and J.W.M. Bertrand** | |
| WP 11 | Implementing statistical process control in industry : the role of statistics and statisticians<br>**W.A.J. Schippers and R.J.M.M. Does** | |
| WP 12 | Delivery performance improvement by controlled work-order release and work center-loading<br>**H.P.G. van Ooijen** | |
| WP 13 | Estimating stock levels in periodic review inventory systems<br>**M.C. van der Heijden** | |
| WP 14 | Near cost-optimal inventory control policies for divergent networks under fill rate constraints<br>**M.C. van der Heijden** | |
| WP 15 | Covering a rectangle with six and seven circles<br>**J.B. Melissen and P.C. Schuur** | |
| WP 16 | Facility layout problemen met vaste oppervlaktes en variabele vorming<br>**P.C. Schuur** | |
| WP 17 | Coordinated planning of preventive maintenance in hierarchical production systems<br>**G. van Dijkhuizen and A. van Harten** | |
| WP 18 | Fuzzy group decision making in a competitive situation<br>**J. Yan, A. van Harten, L. van der Wegen** | |
| WP 19 | Existence of symmetric metaequilibria and their a-priori probability in metagames<br>**J. Yan, A. van Harten, L. van der Wegen** | |
| WP 20 | New circle coverings of an equilateral triangle<br>**J.B.M. Melissen, P.C. Schuur, A. Heppes, Zs. Gaspar en T. Tarnai** | |
| WP 21 | Evaluation of three control concepts for the use of recipe flexiblity in production planning<br>**W.G.M.M. Rutten and J.W.M. Bertrand** | |

WP 22 A taxonomy of systems of corporate governance
**Jeroen Weimer & Joost C. Pape**

WP 23 Bedrijfskunde en methodologie
**Sander van Triest**

WP 24 On the theoretical relation between operating leverage, earnings variability, and systematic risk
**Sander van Triest & Aswin Bartels**

WP 25 Business care Roteb : recovery strategies for monitors
**Harold Krikke**

WP 26 The organisational and information aspects of the financial logistical management concept in theory and practice
**D. Swagerman, A. Wassenaar**

WP 27 Intercontinental airline flight schedule design
**A. van Harten, P.D. Bootsma**

WP 28 An integrated approach to process control
**Werner A.J. Schippers**

WP 29 Inventory control in multi-echelon divergent systems with random lead times
**M.C. van der Heijden, E.B. Diks en A.G. de Kok**

WP 30 Revolution through electronic purchasing
**J. Telgen**

WP 31 Over uitbesteding van voedingsverzorging in de zorgsector
**A. van Harten, E.J. Braakman, D. Schuilenburg, J.P. Papenhuizen**

WP 32 Dynamic scheduling of batch operations with non-identical machines
**D.J. van der Zee, A. van Harten, P.C. Schuur**

WP 33 Multi-echelon inventory control in divergent systems with shipping frequencies
**M.C. van der Heijden**

WP 34 A new magical lamp to rub? : the multiple constituency approach : a potentially useful framework for research on the organizational effectiveness construct
**J. Weimer, M. van Riemsdijk**

WP 35 Het bepalen van een recovery-strategie voor afgeslankte duurzame producten
**H.R. Krikke**

WP 36 Integral organizational renewal : between structure and uncertainty
**F.M. van Eijnatten, L.A. Fitzgerald**

WP 37 Inter-company supply chain planning : extending the current modeling perspective
**J. Fransoo, M. Wouters**

WP 38 Zorginstellingen hebben behoefte aan hulp bij verbetering van de doelmatigheid
**J. Telgen**

WP 39 Spare parts management at the Royal Netherlands Navy : vari-metric and beyond
**J.W. Rustenburg, G.J.J.A.N. van Houtum, W.H.M. Zijm**

WP 40 Evolution of ERP systems
**J.C. Wortmann**

WP 41 A queuing model for due date control in a multi server repair shop with subcontracting
**J. Keizers, I. Adan, J. van der Wal**

WP 42 Stocking strategy for service parts : a case study
**R. Botter, L. Fortuin**

WP 43 Coordinated production maintenance planning in airline flight schedule development
**G. van Dijkhuizen**

WP 44 Warehouse design and control : framework and literature review
**B. Rouwenhorst, R.J. Mantel, B. Reuter, V. Stockrahm, G.J. van Houtum, and W.H.M. Zijm**

WP 45 On a number theoretic property of optimal maintenance grouping
**A. van Harten, G.C. Dijkhuizen & J.G.M. Mars**

WP 46 Spare parts management for technical systems : resupply of spare parts under limited budgets
**W.D. Rustenburg, G.J. van Houtum, W.H.M. Zijm**