

Periodic review for a perishable item under non stationary stochastic demand

Roberto Rossi*

* *University of Edinburgh Business School, Edinburgh, EH8 9JS UK
(e-mail: roberto.rossi@ed.ac.uk).*

Abstract: We consider the periodic-review, single-location, single-product, production/inventory control problem under non stationary demand and service-level constraints. The product is perishable and has a fixed shelf life. Costs comprise fixed ordering costs and inventory holding costs. For this inventory system we discuss a number of control policies that may be adopted. For one of these policies, we assess the quality of an approximate Constraint Programming (CP) model for computing near optimum policy parameters.

Keywords: Inventory control, production planning and scheduling

1. INTRODUCTION

As pointed out in Entrup (2005), Advanced Planning Systems generally tend to not adequately incorporate shelf life aspects of food in their inventory control facilities.

Inventory problems of perishable products have been discussed extensively in the literature. Nahmias (1982) provides a review of the early literature on ordering policies for perishable inventories between 1960s and 1982. Karaesmen et al. (2011) review the more recent supply chain management literature of perishable products having fixed or random lifetimes. For these problems, the structure of the optimal replenishment policy is typically complex: the replenishment quantity depends on the individual age categories of current inventories and all outstanding orders. For this reason, the authors point out that developing effective heuristic policies is of great practical importance in inventory systems for perishable items.

We consider the periodic-review, single-location, single-product, production / inventory control problem under non stationary demand and service-level constraints. The product is perishable and has a fixed shelf life. Costs comprise fixed ordering costs and inventory holding costs. A similar problem was considered in Minner and Transchel (2010); however the authors adopted the simplifying assumption that fixed ordering costs are negligible. For this inventory system we discuss a number of control policies that may be adopted. For one of these policies, we assess the quality of the approximate Constraint Programming (CP) model proposed in Rossi et al. (2010).

2. PROBLEM DESCRIPTION

We consider a planning horizon of N periods and a demand d_t for each period $t \in \{1, \dots, N\}$, which is a non-negative random variable with known probability density function $g_t(d_t)$. We assume that the demand occurs instantaneously at the beginning of each time period. The demand is non-stationary, that is it can vary from period to period and demands in different periods are assumed to be

independent. Demands occurring when the system is out of stock are back-ordered and satisfied as soon as the next replenishment order arrives. The sellback of excess stock is not allowed. A fixed delivery cost a and a proportional unit cost u are incurred for each order. A replenishment order is assumed to arrive instantaneously at the beginning of each period, before the demand in that period occurs. For ease of exposition, we assume that there is no replenishment lead-time; however, the model can be easily extended to systems with positive and fixed replenishment lead-times. Each item that is delivered by the supplier arrives fresh and expires in exactly $M + 1$ periods; therefore a product age may range from 0 to M . A linear holding cost h is incurred for each unit of product carried in stock from one period to the next. A linear wastage cost w is incurred, at the end of each period, for each unit of product that reached age M . Our aim is to find a replenishment plan that minimizes the expected total cost, which is composed of ordering costs, holding costs, and wastage costs over an N -period planning horizon, while satisfying given service level constraints. As service level constraints, we require that, with a probability of at least α , at the end of each period the net inventory will be non-negative.

The actual sequence of actions is to some extent arbitrary. In what follows, we will assume that at the beginning of a period, the inventory on hand after all the demands from previous periods have been realized is known, for each product age that is available. Since we are assuming complete backlogging, this quantity may be negative. However, note that only fresh products can be backordered, since the supplier only delivers fresh products. On the basis of this information, an ordering decision is made for the current period and the respective order is immediately received. Then the period demand is observed and the stock is reduced according to a “first in first out” (FIFO) issuing policy. If, after the demand has been observed, there are still items of age M in stock, these are disposed at cost w per unit. Finally, holding cost is incurred on the remaining stock that is carried over to the next period.

4. OPTIMAL POLICY

$$\min \int_{d_1} \dots \int_{d_N} \sum_{t=1}^N \left(a\delta_t + uQ_t + \max(h \sum_{i=0}^{M-1} I_t^i, 0) + wI_t^M \right) \quad (1)$$

$g_1(d_1) \dots g_N(d_N) dd_1 \dots dd_N$

subject to

$$\delta_t = \begin{cases} 1, & \text{if } Q_t > 0 \\ 0, & \text{otherwise} \end{cases} \quad t = 1, \dots, N \quad (2)$$

$$\sum_{i=0}^M I_t^i + d_t - \sum_{i=0}^{M-1} I_{t-1}^i = Q_t \quad t = 1, \dots, N \quad (3)$$

$$I_t^i = \max \left(I_{t-1}^{i-1} - \max \left(d_t - \sum_{j=i}^{M-1} I_{t-1}^j, 0 \right), 0 \right) \quad \begin{matrix} t = 1, \dots, N \\ i = 1, \dots, M \end{matrix} \quad (4)$$

$$\Pr \left\{ \sum_{i=0}^M I_t^i \geq 0 \right\} \geq \alpha \quad t = 1, \dots, N \quad (5)$$

$$I_0^i = 0 \quad i = 1, \dots, M \quad (6)$$

$$I_t^i \geq 0 \quad \begin{matrix} t = 1, \dots, N \\ i = 1, \dots, M \end{matrix} \quad (7)$$

$$I_t^0 \in \mathbb{R} \quad t = 1, \dots, N \quad (8)$$

$$\delta_t \in \{0, 1\} \quad t = 1, \dots, N \quad (9)$$

$$Q_t \geq 0 \quad t = 1, \dots, N. \quad (10)$$

Fig. 1. Stochastic programming model

3. STOCHASTIC PROGRAMMING MODEL

We extend the chance-constrained model originally proposed in Bookbinder and Tan (1988) and relax the assumption that items can be held in stock for a potentially unlimited amount of time. The chance-constrained model is shown in Fig. 1. Our extended model tracks inventory of different ages via dedicated random variables and issuing policy constraints. More specifically, I_0^i denotes the initial inventory of age $i \in \{0, \dots, M-1\}$, without loss of generality here assumed to be zero (Const. 6). I_t^i denotes inventory level of age $i \in \{1, \dots, M\}$ at the end of period $t \in \{1, \dots, N\}$, a random variable that takes non-negative values, since non-fresh products cannot be ordered (Const. 7). I_t^0 denotes fresh items in stock at the end of period $t \in \{1, \dots, N\}$, a random variable that takes real values, since fresh products can be back-ordered (Const. 8). δ_t is a binary decision variables that is set to 1 if and only if an inventory review is scheduled in period $t \in \{1, \dots, N\}$ (Const. 9). Q_t denotes the non-negative order quantity in period $t \in \{1, \dots, N\}$ (Const. 10); in other words, items in stock cannot be returned to the supplier. The objective function (1) minimizes the expected total cost, which is composed by the fixed ordering/setup cost a for each order placed, the unit ordering cost for each item ordered, the unit holding cost h for products in stock carried from one period to the next, the unit wastage cost w for items that expire. Const. 2 states that a replenishment is scheduled, and the respective fixed cost is incurred in the objective function, if the order quantity is positive at a given period $t \in \{1, \dots, N\}$. Const. 3 is the inventory conservation constraint, which relates the inventory level at the end of period t , the order quantity Q_t in period t , the realized demand d_t in period t , and the inventory carried over from period $t-1$ to period t . Const. 4 implements a FIFO issuing policy and the respective rule for the “consumption” and “aging” of items in stock. Const. 5 enforces the service level in each period.

Deriving the optimal policy for the stochastic program discussed in Section 3 is a non-trivial task. To date, there exists no complete solution method for accomplishing this task for a generic demand distribution. However, when the stochastic demand d_t in period $t = 1, \dots, N$ follows a discrete distribution defined over a finite support, the optimal policy for the stochastic program discussed in Section 3 can be obtained, for small instances, by using a deterministic equivalent scenario based model, see Birge and Louveaux (1997).

Numerical example We consider a planning horizon comprising 4 periods. In each period we observe a random demand that follows a discrete distribution. The probability mass functions for the demand is

$$\begin{aligned} \text{pmf}(d_1) &= \{18(0.5), 26(0.5)\} \\ \text{pmf}(d_2) &= \{52(0.5), 6(0.5)\} \\ \text{pmf}(d_3) &= \{9(0.5), 43(0.5)\} \\ \text{pmf}(d_4) &= \{20(0.5), 11(0.5)\}. \end{aligned}$$

Accordingly, in period 1 we may observe 2 values for the random demand, 18 and 26, each of which occurs with probability 0.5. The complete set of scenarios is presented in Table 1. The fixed delivery cost a is set to 300, the

| Scenario | d_1 | d_2 | d_3 | d_4 | Probability |
|----------|-------|-------|-------|-------|-------------|
| 1 | 18 | 52 | 9 | 20 | 0.0625 |
| 2 | 18 | 52 | 9 | 11 | 0.0625 |
| 3 | 18 | 52 | 43 | 20 | 0.0625 |
| 4 | 18 | 52 | 43 | 11 | 0.0625 |
| 5 | 18 | 6 | 9 | 20 | 0.0625 |
| 6 | 18 | 6 | 9 | 11 | 0.0625 |
| 7 | 18 | 6 | 43 | 20 | 0.0625 |
| 8 | 18 | 6 | 43 | 11 | 0.0625 |
| 9 | 26 | 52 | 9 | 20 | 0.0625 |
| 10 | 26 | 52 | 9 | 11 | 0.0625 |
| 11 | 26 | 52 | 43 | 20 | 0.0625 |
| 12 | 26 | 52 | 43 | 11 | 0.0625 |
| 13 | 26 | 6 | 9 | 20 | 0.0625 |
| 14 | 26 | 6 | 9 | 11 | 0.0625 |
| 15 | 26 | 6 | 43 | 20 | 0.0625 |
| 16 | 26 | 6 | 43 | 11 | 0.0625 |

Table 1. Scenarios

proportional unit cost u to 2, the holding cost h to 1 and the wastage cost w to 4. The shelf life M is set to 2 and the prescribed satisfaction probability α is 0.85. By using a scenario-based deterministic equivalent mixed integer linear programming model we can solve the above instance in reasonable time. The optimal policy is presented in Fig. 2. Black nodes in the policy tree represents orders. The respective order quantity is displayed beside the node. We observe stock-outs in 2 scenarios over 16 (non-stockout probability: 0.875), in both period 3 and 4. Waste is observed at the end of period 3 in scenarios 13, ..., 16.

Unfortunately, an optimal policy is highly unstructured and therefore hardly usable in practice. In what follows we will therefore discuss more structured and usable suboptimal policies, which feature different levels of complexity when it comes to implementation.




Fig. 2. Optimal policy parameters

5. REPLENISHMENT CYCLE POLICY

Different inventory control policies can be implemented for the stochastic program in Section 3. Bookbinder and Tan (1988) define a number of policies that are applicable when items are not perishable. These policies are motivated by practical considerations on inventory handling practices. We developed deterministic equivalent scenario based mixed integer linear programming (MILP) model for computing policy parameters under different strategies. For space reasons, this and the following MILP models are not included in this document.

5.1 “Static uncertainty” strategy

The first strategy introduced by Bookbinder and Tan in Bookbinder and Tan (1988) is the so-called “static-uncertainty” strategy. In this strategy order quantities and review times are fixed once-and-for-all at the beginning of the planning horizon. In practice, this policy may be of interest for practitioners in all those situations in which replenishment periods as well as precise order quantities must be agreed with the customer in advance.

Numerical example. We consider the same instance discussed in Section 4. We solve this instance by using the MILP model developed (Table 2).

Expected total cost: 1065.5

| Period | 1 | 2 | 3 | 4 |
|--------|----|---|----|---|
| Q_t | 78 | 0 | 54 | 0 |

Table 2. Optimal policy parameters under a “static-uncertainty” strategy

In several situations, however, the “static-uncertainty” strategy results not flexible enough. When customer demand is non-stationary and the accuracy of the forecast

is low, Bookbinder and Tan (1988) proposed a more flexible strategy known as “static-dynamic uncertainty”. This strategy features a series of review times, all fixed at the beginning of the planning horizon (i.e., the static aspect of the strategy). This provides an effective means of damping planning instability (deviations in planned orders, also known as nervousness de Kok and Inderfurth (1997)). However, the actual order quantities are determined only after observing the realized demand (i.e., the dynamic aspect of the strategy).

5.2 “Static-dynamic uncertainty” strategy

When items in stock are perishables, the “static-dynamic uncertainty” strategy may be formulated as a stock-age independent or as a stock-age dependent policy.

Stock age independent A stock-age independent “static-dynamic uncertainty” strategy associates with each review period an order-up-to-level. As in the classical “static-dynamic uncertainty” strategy for non-perishable items, the order quantity is computed as the amount of stock required to raise the inventory level up to the order-up-to-level, regardless of the age of products in stock carried over from previous periods. Order-up-to-levels for review periods are set in such a way as to compensate for the realized waste and to ensure the required service level. This strategy may be appealing for practitioners, since it does not require to take into account the different ages of stock on hand. However, it may clearly produce higher waste than a stock-age dependent policy and therefore incur higher expected total costs, since order quantities do not take into account the age, but only the number of items available in stock.

Numerical example We consider the same instance discussed in Section 4. We solve this instance by using the MILP model developed (Table 3).

Expected total cost: 1005.5

| Period | 1 | 2 | 3 | 4 |
|------------|----|---|----|---|
| S_t | 78 | 0 | 66 | 0 |
| δ_t | 1 | 0 | 1 | 0 |

Table 3. Optimal policy parameters under a stock age independent “static-dynamic uncertainty” strategy

Stock age dependent Conversely, a stock-age dependent “static-dynamic uncertainty” strategy does not operate based on order-up-to-levels. For each review period the order quantity is computed as the minimum amount of stock required to guarantee the required service level up until the next review period. This quantity is computed by taking into account the age and the amount of items available in stock. This strategy may guarantee lower waste and expected total cost than a stock-age independent “static-dynamic uncertainty” strategy. However, the computation of the order quantity is more complex than in a stock-age independent strategy. This complicates the implementation of this strategy in practical settings.

Numerical example We consider the same instance discussed in Section 4. We solve this instance by using the MILP model developed (Fig 3).




Fig. 3. Optimal policy parameters under a stock age dependent “static-dynamic uncertainty” strategy

In what follows, we will concentrate on the stock-age dependent “static-dynamic uncertainty” strategy. This strategy is preferred to the others described above because of the flexibility it ensures in order quantity computation. This flexibility ensures lower waste and expected total costs. Furthermore, a decision maker usually relies on advanced ERP systems to automate the computation of actual order quantities. Therefore, although the implementation is slightly more complicated than that of a stock-age independent “static-dynamic uncertainty” strategy, this policy remains viable in practice.

6. A BRUTE FORCE APPROACH FOR THE STOCK AGE DEPENDENT “STATIC-DYNAMIC UNCERTAINTY” STRATEGY

We introduce the notion of replenishment cycle.

Definition 1. A replenishment cycle $T(i, j)$ is the time span between two consecutive replenishments in periods i and $j + 1$, where $i \leq j$.

Definition 2. The cycle opening-inventory-level $R(i, j)$ is minimum opening inventory level in period i that is required in order to satisfy the service level constraints over a replenishment cycle $T(i, j)$.

Finding an optimal non-stationary replenishment cycle policy means finding a set of non-overlapping replenishment cycles and the respective opening-inventory-levels.

Once i and j have been fixed, so to identify a replenishment cycle $T(i, j)$, it is straightforward to compute the cycle opening-inventory-level by using the well-known newsboy formula, see e.g. Axsäter (2010). Consider the cumulative distribution function of $d_i + \dots + d_j$, $G_{d_i + \dots + d_j}(\cdot)$. If we assume that $G_{d_i + \dots + d_j}(\cdot)$ is strictly increasing, $G_{d_i + \dots + d_j}^{-1}(\cdot)$

is uniquely defined. Then it is straightforward to compute the cycle opening inventory level

$$R(i, j) = G_{d_i + \dots + d_j}^{-1}(\alpha) \quad (11)$$

that guarantees a non-stockout probability α in each period $t \in \{i, \dots, j\}$.

When items do not expire, under a replenishment cycle policy, the order quantity at period i is easily computed. The decision maker simply determines the closing-inventory-level at the end of the previous period, I_{i-1} , and she orders a quantity $Q = R(i, j) - I_{i-1}$. Therefore the cycle opening-inventory-level $R(i, j)$ is employed as an order-up-to-level, that is the level up to which stocks have to be raised as a consequence of the order placed. The order quantity is computed via a simple linear relation between the order-up-to-level and the closing-inventory-level at the end of the previous period. Unfortunately, this simple rule cannot be applied if part or all the inventory carried over from period $i - 1$ is going to expire before period j , since the realized waste depends upon the demand distribution. The computation of the order quantity is complex. In fact, the order quantity depends on the order-up-to-level, on the different product ages carried over from previous periods, and on the demand distribution in each period of the cycle. In order to enable the decision maker to implement a non-stationary replenishment cycle policy for perishable items, we introduce an effective strategy for computing the order quantity Q at the beginning of each cycle.

To introduce our strategy, we discuss the problem of computing the minimum order quantity Q that is required to meet prescribed service level constraints during a replenishment cycle $T(i, j)$ when a mix of items with different age categories is already available in the system at the beginning of period i . Let I_{i-1}^m be the available inventory of age m . Consider an array $\mathcal{I} = \{I_{i-1}^0, I_{i-1}^1, \dots, I_{i-1}^M\}$ describing the available inventory at the beginning of period i , before our ordering decision is made. Note that I_{i-1}^0 may be negative in order to keep track of situations in which we start with some backordered demand. For the coming $j - i + 1$ periods, in each period $t \in \{i, \dots, j\}$ we observe a normally distributed demand d_t with probability density function $g(d_t)$, mean μ_t and standard deviation σ_t . There is a service level constraint enforcing a non-stockout probability α in each period t . Consider period t , where $i \leq t \leq j \leq M + 1$, since an item can be used only over $M + 1$ periods. Items are issued according to a FIFO issuing policy. The service level constraint for period t can be written as

$$\Pr\{I_t^0 \geq 0\} \geq \alpha. \quad (12)$$

In other words, only fresh items can be backordered. We now introduce the following stochastic recurrence relation

$$I_t^m = \max(I_{t-1}^{m-1} - \max(d_t - \sum_{k=m}^{M-1} I_{t-1}^k, 0), 0), \quad (13)$$

for $t = i, \dots, j$ and $m = 2, \dots, M$. Furthermore,

$$I_t^0 = \min(0, \sum_{k=0}^{M-1} I_{t-1}^k - d_t), \quad (14)$$

for $t = i + 1, \dots, j$; and

$$I_t^0 = Q + \min(0, \sum_{k=0}^{M-1} I_{t-1}^k - d_t), \quad (15)$$

for $t = i$. We also consider the indicator function

$$f(Q, d_i, d_{i+1}, \dots, d_t) = \begin{cases} 1 & \text{if } I_t^0 \geq 0 \\ 0 & \text{otherwise} \end{cases},$$

where I_t^0 is computed according to Eq. 13, 14 and 15.

Given the array \mathcal{I} , describing the available inventory at the beginning of period i , before our ordering decision is made, and an ordering decision Q , by using the indicator function introduced in concert with Eq. 13, 14 and 15, we express the service level constraint as

$$\int_{d_i} \dots \int_{d_t} f(Q, d_i, \dots, d_t) g(d_i) \dots g(d_t) dd_i \dots dd_t \geq \alpha. \quad (16)$$

The left hand side of Eq. 16 is increasing in Q , therefore the minimum order quantity that satisfies the above relation can be found using a binary search procedure that numerically integrates the expression. Due to the cost structure of the stochastic program in Section 3, it is clear that, for a given replenishment cycle, the minimum Q that satisfies Eq. 16 also minimizes the expected total cost for that cycle. In our numerical experiments, we employ Monte Carlo integration to numerically integrate Eq. 16 with a precision of ± 0.005 at 95% confidence.

It should be noted that, if the net inventory is negative (i.e. we begin the cycle with backorders) Q is equal to the cycle opening-inventory-level plus the items backordered and no search is needed. On the other hand, if the net inventory at the beginning of the cycle is positive we can identify two limit cases: if no item expires during the cycle, Q is equal to the difference between the cycle opening-inventory-level and the inventory carried over from previous periods. Conversely, if all the on-hand inventory expires immediately, Q is simply equal to the cycle opening-inventory-level. These two limit cases provide bounds within which the binary search has to be performed.

Since the computation of the order quantity Q for a given replenishment cycle $T(i, j)$ has been now made explicit. Given a set of replenishment cycles, the decision maker can implement the non-stationary replenishment cycle policy for perishable items by employing, at the beginning of each replenishment cycle, the procedure discussed above for computing the associated minimum order quantity Q . Since unit, holding and wastage costs are all increasing in Q , the minimum feasible value for Q is cost-optimal.

Numerical example For the same instance discussed in Section 4, we consider a replenishment cycle that starts in period 3 and ends in period 4. The initial inventory array is $\mathcal{I} = \{44, 2, 0\}$. The procedure discussed prescribes an optimal order quantity $Q = 17$.

If the planning horizon comprises a limited number N of periods (i.e. up to 20 periods), it is possible to find a set of non-overlapping replenishment cycles that minimize the expected total cost under a stock age dependent “static-dynamic uncertainty” strategy, by using a “brute force” approach. In other words, we can try all the possible 2^N combinations of review periods and then estimate by simulation and confidence interval analysis the expected total cost of the heuristic stock age dependent “static-dynamic uncertainty” strategy discussed.

Numerical example We consider the same instance discussed in Section 4. We solve this instance by using the “brute force” approach. The resulting policy places orders in period 1 and 3, the respective order quantities can be computed at the beginning of a given replenishment period via the binary search approach introduced above once demand in previous periods has been observed. The expected total cost of this strategy is 1006, about 3% costlier than the optimal stock age dependent “static-dynamic uncertainty” strategy (i.e. 973.5).

However, in order to plan weekly production for a year (i.e. $N = 36$ weeks) this approach is not viable due to the large number of review period combinations that have to be assessed by simulation. It is therefore essential to develop heuristic approaches to the stock age dependent “static-dynamic uncertainty” strategy. These strategies should compute a near-optimal set of non-overlapping replenishment cycles. Furthermore, they should also provide a good approximation of the expected total cost a decision maker expects to face by employing such a set of review periods. The quality of these heuristics can be assessed, for small instances, against the solution produced by the “brute force” approach just discussed.

6.1 A CP model for computing re-order points

Following a modeling strategy that resembles the one discussed in Tarim and Kingsman (2004) and Tarim and Smith (2008), Rossi et al. (2010) discussed a heuristic CP model for solving the stochastic program in Section 3 under a stock age dependent “static-dynamic uncertainty” strategy. The model provides a set of review periods and an estimation of the total cost the decision maker is expected face if inventory reviews are planned according to this plan. As discussed in Section 6, once a set of review periods is known, we can dynamically integrate Eq. 16 to determine the respective order quantities according to a stock age dependent “static-dynamic uncertainty” strategy.

Numerical example We consider the same instance discussed in Section 4. We solve this instance by using the CP model discussed in Rossi et al. (2010). The CP model is solved by using the normally distributed demands in Table 4 from which the probability mass functions in Section 4 were sampled. The resulting policy correctly suggests to

| Demand | d_1 | d_2 | d_3 | d_4 |
|----------|-------|-------|-------|-------|
| μ | 22 | 29 | 26 | 16 |
| σ | 4 | 23 | 17 | 5 |

Table 4. Normally distributed demands

place orders in period 1 and 3. The estimated expected total cost of this strategy, i.e. objective function of the CP model at optimality, is 951, that is 5% less than the actual cost (i.e. 1006) we observe when we adopt this replenishment plan and we compute order quantities by using the strategy discussed in Section 6.

7. COMPUTATIONAL EXPERIMENTS

We consider a demand that is normally distributed in each period of the planning horizon. In the following patterns

Pattern

- 1 → 8, 9.5, 2, 9, 8, 1.5, 6.5, 8, 9, 3, 1.5, 6
- 2 → 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6
- 3 → 6, 7.3, 8.5, 9, 8.5, 7.3, 6, 4.7, 3.5, 3, 3.5, 4.7
- 4 → 1.5, 2, 3.5, 6, 8, 8.5, 9, 10.5, 9.5, 6.5, 5, 2
- 5 → 19, 9.5, 0.4, 0.8, 0.3, 1.5, 8, 9.5, 11, 3.5, 1.5, 7

figures represent expected demand in '00 units for each of the twelve periods in the planning horizon. The five patterns considered are derived from Berry (1972) and represent erratic (1), stationary (2), seasonal (3), life cycle (4), highly erratic (5) demand, respectively. The model parameters are $N = 12$, $M = 2$ (shelf life of 3 periods), and $a = 3000$, $h = 1$, $u = 2$. The remaining parameters range in the following sets, $\alpha = \{0.90, 0.95, 0.98\}$, $w = \{0, 2, 4\}$, and $\sigma_{d_i} = \{1/3, 1/4, 1/10\}$, where σ_{d_i} denotes the standard deviation of the demand in period $i = 1, \dots, N$.

We compared the policies produced by the CP model against those obtained via the brute force approach discussed in Section 6. Expected total cost of these policies has been estimated by simulation. In the estimation of the expected total cost we allowed a maximum error of $\pm 1\%$ at 95% confidence. In Fig. 4 we reported results for our analysis by using box plots. The average cost difference observed is 0.21%; most of the dispersion lies within $\pm 1\%$. This shows that the CP model generates near optimal policies. We also investigated how the cost predicted by




Fig. 4. Cost difference (in percentage of the optimum policy cost) between CP policies and optimum policies obtained with the approach discussed in Section 6.

the CP model approximates the actual cost of the policy generated. The cost prediction errors are reported in Fig. 5. From these results it is clear that the CP model tends to underestimate costs. However, it is apparent that this underestimation is very low, in fact, on average, the actual cost is underestimated by -0.68%. Most of the dispersion lies within 0 and 1%, with the exception of the highly erratic pattern, for which the under estimation is slightly higher. This demonstrates that the CP model not only generates near optimal policies, but also provides a good approximation of the cost of these policies.




Fig. 5. Cost prediction error (in percentage of the actual cost incurred by the policy) for the CP model.

REFERENCES

- Axsäter, S. (2010). *Inventory Control (International Series in Operations Research & Management Science)*. Springer, 2nd ed. edition.
- Berry, W.L. (1972). Lot sizing procedures for requirements planning systems: A framework for analysis. *Production and Inventory Management Journal*, 13(2), 19–34.
- Birge, J.R. and Louveaux, F. (1997). *Introduction to Stochastic Programming (Springer Series in Operations Research and Financial Engineering)*. Springer.
- Bookbinder, J.H. and Tan, J.Y. (1988). Strategies for the Probabilistic Lot-Sizing Problem with Service-Level Constraints. *Management Science*, 34(9), 1096–1108.
- de Kok, T. and Inderfurth, K. (1997). Nervousness in inventory management: Comparison of basic control rules. *European Journal of Operational Research*, 103(1), 55–82.
- Entrup, M.L. (2005). *Advanced Planning in Fresh Food Industries: Integrating Shelf Life into Production Planning (Contributions to Management Science)*. Physica.
- Karaesmen, I., SchellerWolf, A., and Deniz, B. (2011). Managing perishable and aging inventories: Review and future research directions. In K.G. Kempf, P. Keskinocak, and R. Uzsoy (eds.), *Planning Production and Inventories in the Extended Enterprise*, volume 151 of *International Series in Operations Research & Management Science*, 393–436. Springer US.
- Minner, S. and Transchel, S. (2010). Periodic review inventory-control for perishable products under service-level constraints. *OR Spectrum*, 32(4), 979–996.
- Nahmias, S. (1982). Perishable inventory theory: A review. *Operations Research*, 30(4), 680–708.
- Rossi, R., Tarim, S.A., Hnich, B., and Prestwich, S. (2010). Computing replenishment cycle policy parameters for a perishable item. In *Proceedings of the 24th MEC-EurOPT-2010, Izmir, Turkey*, 241–247.
- Tarim, S. and Kingsman, B.G. (2004). The stochastic dynamic production/inventory lot-sizing problem with service-level constraints. *International Journal of Production Economics*, 88(1), 105–119.
- Tarim, S. and Smith, B. (2008). Constraint programming for computing non-stationary (R,S) inventory policies. *European Journal of Operational Research*, 189(3), 1004–1021.