Periodic review for a perishable item under nonstationary stochastic demand

Dr. Roberto Rossi

The University of Edinburgh Business School, The University of Edinburgh, UK
roberto.rossi@ed.ac.uk

Friday, June the 21th, 2013

Motivation (1)

Nahmias (1982): review of the early literature on ordering policies for perishable inventories between 1960s and 1982.

Karaesmen et al. (2011), Bakker et al. (2013): review the more recent supply chain management literature of perishable products having fixed or random lifetimes.

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

Motivation (2)

In stochastic inventory control for perishable items 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, developing effective heuristic policies is of great practical importance in inventory systems for perishable items (Karaesmen et al., 2011).

Outline of the work

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).

Problem description

We consider a planning horizon of N periods and a demand d_t for each period $t \in \{1, \ldots, N\}$, 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.

Problem description

Sequence of events

The actual sequence of actions is to some extend 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.

Stochastic programming model

$$\min \int_{d_1} \cdots \int_{d_N} \sum_{t=1}^{N} \left(a\delta_t + uQ_t + \max(h \sum_{i=0}^{M-1} I_t^i, 0) + wI_t^M \right)$$
$$g_1(d_1) \ldots g_N(d_N) \mathrm{d}d_1 \ldots \mathrm{d}d_N \qquad (1)$$

subject to

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

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

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

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

$$I_0^i = 0 \qquad i = 1, \ldots, M \qquad (6)$$

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

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

$$\delta_t \in \{0, 1\} \qquad t = 1, \ldots, N \qquad (9)$$

$$Q_t \geq 0 \qquad t = 1, \ldots, N \qquad (10)$$

Optimal policy

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 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 (1)

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

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

The fixed delivery cost a is set to 300, the 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.

Numerical example (2)

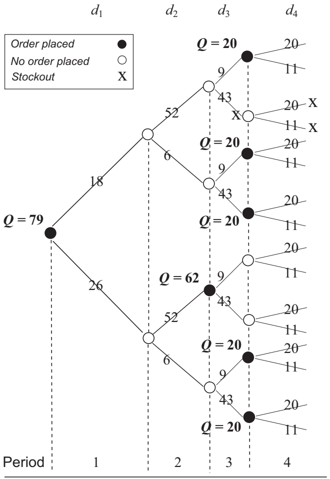

Expected total cost: 918.5

Numerical example (2)

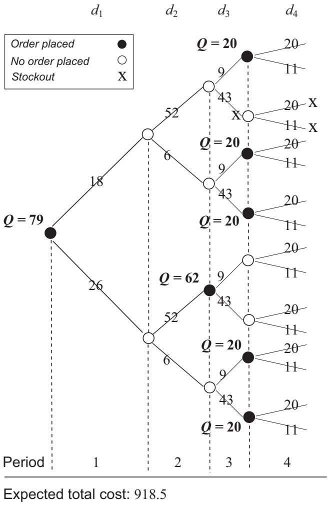

Unfortunately, an optimal policy is highly unstructured and therefore hardly usable in practice.

Heuristic policies

Static-uncertainty

Introduced by Bookbinder and Tan (1988) the so-called "static-uncertainty" strategy fixes order quantities and review times 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

Expected total cost: 1065.5

Period	1	2	3	4
Q_t	78	0	54	0

Table: Optimal policy parameters under a "static-uncertainty" strategy

Heuristic policies

Static-dynamic uncertainty

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). However, the actual order quantities are determined only after observing the realized demand (i.e., the dynamic aspect of the 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.

Static-dynamic uncertainty

stock-age independent policy

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.

Static-dynamic uncertainty

stock-age independent policy

Numerical example

Expected total cost: 1005.5

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

Table: Optimal policy parameters under a stock age independent "static-dynamic uncertainty" strategy

Static-dynamic uncertainty

stock-age dependent policy

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 that in a stock-age independent strategy. This complicates the implementation of this strategy in practical settings.

Static-dynamic uncertainty

stock-age dependent policy

Numerical example

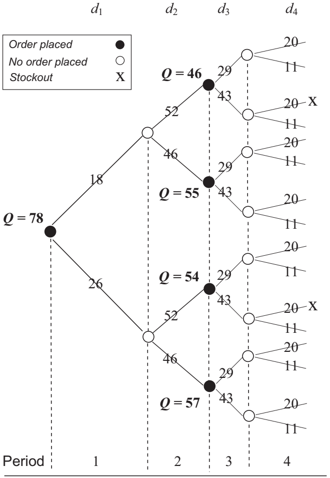

Expected total cost: 951.5

Static-dynamic uncertainty

Computing policy parameters: a brute force approach

Consider the problem of computing the minimum order quantity Q that is required to meet prescribed service level constraints during a replenishment cycle over periods $i, \ldots, j$, 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 $\mathcal{I} = \{I_{i-1}^0, I_{i-1}^1, \ldots, I_{i-1}^M\}$ 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.

Static-dynamic uncertainty

Computing policy parameters: a brute force approach

The service level constraint for period t can be written as

$$\Pr\{I_t^0 \geq 0\} \geq \alpha. \tag{11}$$

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

$$I_t^m = \max(I_{t-1}^{m-1} - \max(d_t - \sum_{k=m}^{M-1} I_{t-1}^k, 0), 0), \tag{12}$$

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

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

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

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

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

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

Static-dynamic uncertainty

Computing policy parameters: a brute force approach

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, we express the service level constraint as

$$\int_{d_i} \dots \int_{d_t} f(Q, d_i, \dots, d_t) g(d_i) \dots g(d_t) \mathrm{d}d_i \dots \mathrm{d}d_t \geq \alpha. \tag{15}$$

The left hand side of Eq. 15 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, for a given replenishment cycle, the minimum Q that satisfies Eq. 15 also minimizes the expected total cost for that cycle.

In our numerical experiments, we employ Monte Carlo integration to numerically integrate Eq. 15 with a precision of ± 0.005 at 95% confidence.

Static-dynamic uncertainty

Computing policy parameters: a brute force approach

Numerical example (1) For the same instance discussed before, 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 $Q = 17$.

Numerical example (2) We consider the same instance discussed before. 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.

Static-dynamic uncertainty

Computing policy parameters: a CP model

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 under a stock age dependent "static-dynamic uncertainty" strategy.

Numerical example The CP model is solved by using the normally distributed demands in Table 3 from which the probability mass functions used in previous examples were sampled.

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

Table: Normally distributed demands

The resulting policy correctly suggests to 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 brute force strategy.

Computational experience

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

Pattern	
1	$\rightarrow 8, 9.5, 2, 9, 8, 1.5, 6.5, 8, 9, 3, 1.5, 6$
2	$\rightarrow 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6, 6$
3	$\rightarrow 6, 7.3, 8.5, 9, 8.5, 7.3, 6, 4.7, 3.5, 3, 3.5, 4.7$
4	$\rightarrow 1.5, 2, 3.5, 6, 8, 8.5, 9, 10.5, 9.5, 6.5, 5, 2$
5	$\rightarrow 19, 9.5, 0.4, 0.8, 0.3, 1.5, 8, 9.5, 11, 3.5, 1.5, 7$

figures represent expected demand in '00 units for each of the twelve periods in the planning horizon. The five patterns considered 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 $i = 1, \ldots, N$.

We compared the policies produced by the CP model against those obtained via the brute force approach. 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.

Computational experience

Goodness of CP model policies

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.

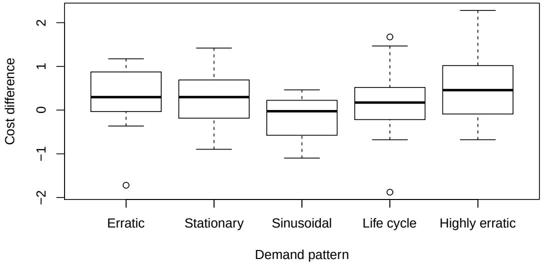

Figure: Cost difference (in percentage of the optimum policy cost) between CP policies and optimum policies obtained with the brute force approach.

Computational experience

Cost approximation of the CP model

We also investigated how the cost predicted by the CP model approximates the actual cost of the policy generated. 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%.

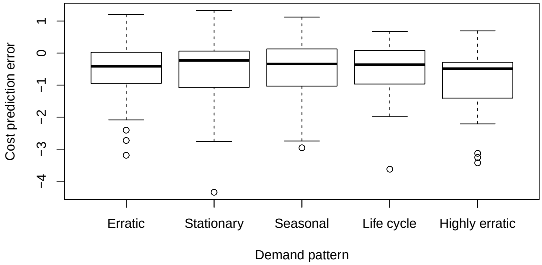

Figure: Cost prediction error (in percentage of the actual cost incurred by the policy) for the CP model.

Conclusions

We considered 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 discussed a number of control policies that may be adopted.

For one of these policies, i.e. stock-age dependent static-dynamic uncertainty, we assess the quality of an approximate Constraint Programming (CP) model for computing near optimum policy parameters.

Our results suggest that the CP model not only generates near optimal policies, but also provides a good approximation of the cost of these policies.