A sample based method for perishable good inventory control with a service level constraint

Eligius M.T. Hendrix[1], Karin G.J. Pauls-Worm[2], Roberto Rossi[3], Alejandro G. Alcoba[1], and Rene Haijema[2]

[1] Computer Architecture, Universidad de Málaga
`{eligius,agutierrez}@uma.es`
[2] Operations Research and Logistics, Wageningen University
`{karin.pauls,rene.haijema}@wur.nl`
[3] Business school, Edinburgh University
`roberto.rossi@ed.ac.uk`

Abstract. This paper studies the computation of order up to levels for a stochastic programming inventory problem of a perishable product. Finding a solution is a challenge as the problem enhances a perishable product, fixed ordering cost and non-stationary stochastic demand with a service level constraint. An earlier study [6] derived order-up-to values via an MILP approximation. We consider a computational method based on the so-called Smoothed Monte Carlo method using sampled demand to optimize values. The resulting MINLP approach uses enumeration, bounding and iterative nonlinear optimization.

Keywords: Inventory control, Perishable products, MINLP, chance constraint, Monte Carlo

1 Introduction

The basis of our study is a Stochastic programming (SP) model published in [6] for a practical production planning problem over a finite horizon of T periods of a perishable product with a fixed shelf life of J periods. The static-dynamic YS policy in [6] provides the decision maker with a list of order timings Y and order-up-to levels S. Their approach to generate the values for the policy is using an approximation based on MILP. The approximation has problems with fulfilling the so-called service level constraints for some of the tested instances. The question in the current paper is how to generate optimal timing and order-up-to values based on Monte Carlo samples of demand; we call this a sample based approach.

The considered problem has uncertain and non-stationary demand such that one produces to stock. Any unmet demand is backlogged. Mathematically this means that negative inventory may occur. To keep waste due to out-dating low, one issues the oldest product first, i.e. FIFO issuance. The model assumes a zero lead time. The investigated model aims to guarantee the customer that the probability of not being out-of-stock is higher than a required service level

α in every period $t \in \{1, 2, ..., T\}$. I.e. the probability of being out of stock is small; less than $1 - \alpha$. The latter implies being mathematically confronted with a chance constraint. This leads to challenging optimization problems, e.g. [2].

Due to dealing with a perishable product, the inventory consists of items of different ages $j = 1, \ldots, J - 1$. Depending on its value at moment t, an order policy should advice the decision maker on the order quantity Q_t. The static-dynamic YS policy provides a list of order moments $Y \in \{0, 1\}^T$ called an order timing vector, or alternatively it provides at each order the so-called replenishment cycle R_t up to the next order takes place. Given the order timing, we derive theoretical results about the order quantity for the given service level constraints. Besides it provides a list with order-up-to levels S_t. The values generated by the presented MILP model of [6] unfortunately do not exactly fulfil the service level (chance) constraints for all instances. Therefore, our first research question is how to generate values for S_t such that the chance constraints are fulfilled for all instances and expected costs are minimized. Following the concept of Monte Carlo estimation, we present a sample based model to generate the values S_t. It is known from literature that the corresponding sample based model called MC-MILP for most instances cannot be solved in reasonable time, e.g. [7]. Therefore, we investigate the possibility to use an equivalent MINLP model based on the Smoothed Monte Carlo method, see [3]. A specific algorithm is designed that uses enumeration and bounding for the integer part Y of the problem leaving us with iteratively solving an NLP problem in the continuous variables S.

We present a more flexible variant of the YS policy that we call YQ(X) policy. It also fixes the best order moments Y but in the order quantity Q_t takes the distribution X of the age of items in stock into account. The second question is how to determine the order quantity. The YQ(X) policy is more difficult to implement for decision support in practice than the YS policy. Our third question is for which cases the YS policy is doing much worse than the YQ(X) policy. We measure how well required service levels are met and estimate the expected cost of all generated policies for the instances provided by [6].

This paper is organised as follows. Section 2 describes the underlying SP model with chance constraints. Section 3 focuses on the properties of the optimal order quantities considering the corresponding basic order-up-to levels and the concept of approximating the chance constraints by samples. Section 4 describes the MC-MILP and MINLP models that can be used to generate values for the quantities S_t for the YS policy. Section 5 provides the elaboration of a sample based algorithm to optimize the YQ(X) policy. In Section 6, we measure numerically for which characteristics of the instances the YS policy provides good enough performance compared to the YQ(X) policy. Section 7 summarizes our findings.

2 Stochastic Programming Model

The stochastic demand implies that the model has random (denoted in bold) inventory variables $\boldsymbol{I}_{jt}$ apart from the initial fixed levels I_{j0}. If the order decision Q_t depends on the inventory levels at the beginning of the period, then Q_t is also a random variable. In the notation, $P(.)$ denotes a probability to express the chance constraints and $E(.)$ is the expected value operator for the expected costs. Moreover, we use $x^+ = max\{x, 0\}$. A formal description of the SP model presented in [6] is given.

Indices
    t period index, $t = 1, \ldots, T$, with T the time horizon
    j age index, $j = 1, \ldots, J$, with J the fixed shelf life
Data
    $\boldsymbol{d}_t$ Normally distributed demand with expectation $\mu_t > 0$ and variance $(cv \times \mu_t)^2$
        where cv is a given coefficient of variation.
    k fixed ordering cost, $k > 0$
    c unit procurement cost, $c > 0$
    h unit inventory cost, $h > 0$
    w unit disposal cost, is negative when having a salvage value, $w > -c$
    α required service level, $0 < \alpha < 1$
Variables
    $Q_t \geq 0$ ordered and delivered quantity at the beginning of period t
    $\boldsymbol{I}_{jt}$ Inventory of age j at end of period t, initial inventory fixed $I_{j0} = 0$,
        $\boldsymbol{I}_{1t} \in \mathbb{R}$, $\boldsymbol{I}_{jt} \geq 0$ for $j = 2, \ldots, J$.

The total expected costs over the finite horizon is to be minimized.

$$E\left(\sum_{t=1}^{T}\left(g(Q_t) + h\sum_{j=1}^{J-1}\boldsymbol{I}_{jt}^+ + w\boldsymbol{I}_{Jt}\right)\right) = \sum_{t=1}^{T}E\left(g(Q_t) + h\sum_{j=1}^{J-1}\boldsymbol{I}_{jt}^+ + w\boldsymbol{I}_{Jt}\right),$$
(1)

where procurement cost is given by the function

$$g(x) = k + cx, \text{ if } x > 0, \text{ and } g(0) = 0.$$
(2)

The chance constraint expressing the required service level is

$$P(\boldsymbol{I}_{1t} \geq 0) \geq \alpha, \ t = 1, \ldots, T.$$
(3)

Due to FIFO issuing, if the freshest inventory is negative, no older items are in stock. The inventory dynamics of items of different ages is described by

$$\boldsymbol{I}_{1t} = Q_t - (\boldsymbol{d}_t - \sum_{j=1}^{J-1}\boldsymbol{I}_{j,t-1})^+, \ t = 1, \ldots, T$$
(4)

for the freshest inventory and

$$\boldsymbol{I}_{jt} = \left(\boldsymbol{I}_{j-1,t-1} - (\boldsymbol{d}_t - \sum_{i=j}^{J-1}\boldsymbol{I}_{i,t-1})^+\right)^+, \ t = 1, \ldots, T, j = 2, \ldots, J$$
(5)

for older items. These dynamic equations describe the FIFO issuing policy and imply that I_{1t} is a free variable, whereas I_{jt} is nonnegative for the older vintages $j = 2, \ldots, J$. Notice that the oldest inventory I_{Jt} perishes and becomes waste. Let

$$X = (I_{1,t-1}, \ldots, I_{J-1,t-1}) \tag{6}$$

be the inventory with its age distribution at the beginning of the period. An order policy is a set of rules $Q_t : \mathbb{R}^{J-1} \to \mathbb{R}, t = 1, \ldots, T$ which given the inventory state X at the beginning of the period specifies the amount to be ordered by function $Q_t(X)$.

For a nonperishable product, a common way to deal with the planning is to define so-called order-up-to levels S_t, e.g. [9]. The decision maker replenishes at order moment t the inventory up to a level S_t. This means a vector S is provided to the decision maker and the order quantity is defined by

$$Q_t(X) = (S_t - \sum_{j=1}^{J-1} X_j)^+, \ t = 1, \ldots, T. \tag{7}$$

Considering a fixed timing vector Y combined with order-up-to levels S_t is called an YS policy in [6]. The question is how to generate good values for S_t in case we are dealing with a perishable product and part of the inventory will become waste. We deal with this question in Sections 3 and 4.

Order policies that take the age distribution into account have been studied for stationary demand, see [10]. However, following this concept in a replenishment cycle environment with non-stationary demand requires a time dependent rule $Q_t(X)$. Considering a fixed timing vector Y combined with age dependent order quantities $Q_t(X)$ is called an YQ(X) policy here. The decision maker needs an information system that advices on the order quantity $Q_t(X)$; this is practically more complicated than the YS policy. In Section 5, we investigate how order quantities can be derived using sampled demand. As the YQ(X) policy is wider than the YS policy, it should provide lower expected cost than the case where the age distribution is not taken into account.

First we study the implications of order quantities to fulfil the service level constraint when the order moments are fixed beforehand by vector Y.

3   Replenishment cycles, basic order-up-to levels and estimating the service level

In the policies under consideration, the decision maker is provided an order timing vector Y, i.e. $Y_t = 0 \Rightarrow Q_t = 0$. We first focus on the concept of replenishment cycles in Section 3.1 and determine in which cases the so-called basic order-up-to level is the optimal quantity in Section 3.2. For the other order moments, we study the mathematical implications of estimating the service level by a Monte Carlo sampling approach in Section 3.3.

3.1   Replenishment cycles and limits on timing vector Y

Literature on inventory control e.g. [9] applies the concept of a replenishment cycle, i.e. the length R of the period for which the order of size Q_t is meant. For stationary demand, the replenishment cycle is fixed, but for non-stationary demand the optimal replenishment cycle R_t may depend on order moment t. In our case, replenishment cycle R_t for order moment t with $Y_t = 1$ is the number of periods such that the relations $Y_t = Y_{t+R_t} = 1$ and $Y_p = 0, p = t+1, .., t+R_t - 1$ hold. Notice that for the perishable case with shelf life J, practically the replenishment cycle can not be larger than the shelf life J; so $R_t \leq J$.

Lemma 1. *Let Y be an order timing vector of the SP model, i.e. $Y_t = 0 \Rightarrow Q_t = 0$. Y provides an infeasible solution of the SP model, if it contains more than $J - 1$ consecutive zeros.*

This means that a feasible order timing vector Y does not contain a consecutive series with more than $J - 1$ zeros.

Let F_T be the set of all feasible order timing vectors Y of length T. The number of elements $|F_T|$ of the set F_T of feasible order timings of T periods and a shelf life of $J + 1 < T$ follows the recursive rule $|F_{T+1}| = 2|F_T| - |F_{T-J}|$ with the initial terms $|F_t| = 2^{t-1}$ for $t < J + 1$ and $|F_{J+1}| = 2^J - 1$ as shown in [1]. F_T is exponential in the horizon T. However, in a practical situation, as explained in [6], one cannot plan far ahead, because the forecasts of demand on which he distribution of d_t is based, is not known. Therefore, we will focus on the provided instances in that paper with $T = 12$.

3.2   Optimal order quantities and basic order-up-to level

A concept that uses the replenishment cycle is that of the basic order-up-to level. In this context, we define the basic order-up-to level $\hat{S}_{R,t}$ as the inventory that should be available at the beginning of period t to cover demand of R periods.

Definition 1. *Let $\boldsymbol{d}_t + .. + \boldsymbol{d}_{t+R-1}$ be the stochastic demand during the replenishment cycle of length R with cumulative distribution function (cdf) $G_{R,t}$. The basic order-up-to level $\hat{S}_{R,t}$ with probability α to fulfil demand is defined by $G(\hat{S}_{R,t}) = \alpha$ such that $\hat{S}_{R,t} = G_{R,t}^{-1}(\alpha)$.*

For the SP problem describing the inventory development of a perishable product, to have $\hat{S}_{R,t}$ in stock at the beginning of period t, i.e. $\sum_j X_t + Q_t \geq \hat{S}_{R,t}$, may not be sufficient, because products in stock may perish during the replenishment cycle. However, for some of the order moments it may be sufficient and also defines the optimum order quantity. First consider an order moment following a replenishment cycle of the length of the shelf life J.

Lemma 2. *Let Y be an order timing vector of the SP model with corresponding cycle R_t and X defined by (6). For an order moment t having $Y_{t-J} = 1, R_{t-J} = J$, the optimal order quantity is $Q_t = \hat{S}_{R_t,t}$.*

Proof. The cost minimisation aims at a value of Q_t as low as possible. Constraints (4) and (5) define that after J periods no (non-perished) inventory is left over from order Q_{t-J}, so $X_{tj} = 0$ for $j = 1, \ldots, J-1$. To fulfil chance constraint (3), the order quantity should fulfil $Q_t \geq \hat{S}_{R_t,t}$. Minimising its value implies $Q_t = \hat{S}_{R_t,t}$.

$\square$

Second, we may have replenishment cycles of just one period where during the cycle, no products can perish.

Lemma 3. *Let Y be an order timing vector of the SP model and X defined by (6). For an order moment t having $Y_{t+1} = Y_t = 1$ the optimal order quantity is $Q_t = \hat{S}_{1,t} - \sum_j X_j$.*

Proof. For one period demand, chance constraint (3) translates to $P(Q_t + \sum_j X_j \geq \boldsymbol{d_t}) = \alpha \Rightarrow Q_t \geq \hat{S}_{1,t} - \sum_j X_j$. Minimising its value implies $Q_t = \hat{S}_{1,t} - \sum_j X_j$.

$\square$

Independently of the order timing, the best order quantity at a negative stock level always has an order-up-to character.

Lemma 4. *Let Y be an order timing vector of the SP model with corresponding replenishment cycles R_t and X defined by (6). If for $Y_t = 1$, $X_1 \leq 0$ the optimal order quantity is $Q_t = \hat{S}_{R,t} - X_1$.*

Proof. No old stock is available that can perish during the replenishment cycle. The chance constraint (3) translates to $Q_t + X_1 \geq \hat{S}_{R_t,t}$. Minimising the value of Q_t implies $Q_t = \hat{S}_{R_t,t} - X_1$.

$\square$

Given an order timing vector Y, the theoretical properties give us a hand to determine the optimal order quantities for some of the order moments. The question now is how to deal with the chance constraint and the order quantities for the other order moments. A usual way to deal with that is using samples of the demand series.

3.3   Monte Carlo estimation of the service level

Let for an order period t, $\boldsymbol{d}$ be the stochastic demand vector $(\boldsymbol{d_t}, \ldots, \boldsymbol{d_{t+R-1}})$ during replenishment cycle R, X the starting inventory and Q the order quantity. Let $f(Q, X, d) = I_{1,t+R-1}$ define the end inventory of items with age one period given a realisation d of $\boldsymbol{d}$ following the inventory dynamics with possible perishing according to (4) and (5). Consider the indicator function $\delta : \mathbb{R} \times \mathbb{R}^R \rightarrow \{0, 1\}$

$$\delta(Q, d) = \begin{cases} 1 & \text{if } f(Q, X, d) \geq 0 \\ 0 & \text{otherwise} \end{cases} \tag{8}$$

translating the service level in constraint (3) for period $t + R - 1$ to

$$a(Q) = P(I_{1,t+R-1} \geq 0) = E_d \delta(Q, d). \tag{9}$$

The generic concept of his translations is given in handbooks like [4]. Using sample paths $d_1, \ldots, d_N$ of $\boldsymbol{d}$ to estimate a probability like (9) was called by von Neumann the Monte Carlo method. The idea is that given N sample paths $d_1, \ldots, d_N$ of $\boldsymbol{d}$, the probability (service level) (9) is estimated by

$$\hat{a}(Q) = \frac{1}{N} \sum_{r=1}^{N} \delta(Q, d_r). \tag{10}$$

As is know from handbooks on statistics (e.g.[5]), considering a set of independent random samples $\boldsymbol{d}_r$ provides the unbiased estimator (10) of $a(Q)$ with standard deviation

$$\sigma(\hat{a}(Q)) = \sqrt{\frac{1}{N}(a(Q) - a(Q)^2)}. \tag{11}$$

The latter is of importance in Monte Carlo approaches to set the number of samples for a desired probabilistic accuracy. Following the usual idea that the binomial distribution is practically normal for a large number of samples, a rule of thumb is to have an accuracy of 2σ. For the particular application aiming at $\alpha = 0.90, 0.95, 0.98$, a sample size of $N = 5000$ gives a rule of thumb accuracy of about 0.005 of the estimator $\hat{a}(Q)$.

The next question is how to use the theoretical findings of Section 3.2 and the estimation method of Section 3.3 in order to find policies where the order timing Y is provided to the decision maker. This means, we should find the best order timing and a way to deal with the order quantity.

4   fixed timing Y, order-up-to level S_t

In the YS policy, the decision maker is provided a list of order-up-to levels S_t for each order moment and orders according to (7). Lemmas 2 and 3 are helpful to define the order-up-to level $S_t = \hat{S}_{R,t}$ for specific moments. Sample based estimation can be used for the service level in each period, for this policy defined as a vector $a(S) = (a_1(S), \ldots, a_T(S))$.

Specifically, for the YS policy, one can write the problem of finding the (discrete) timing Y and (continuous) order-up-to levels S as a Monte Carlo based Mixed Integer Linear Programming (MC-MILP) model. Such a model is given in Section 4.1. When using samples to measure a probability (service level) as function of a continuous variable, the resulting function is piecewise constant in the continuous variable, here S. To be more precise, the implicit equivalent of estimator (10) is a function $\hat{a}_t(S) : \mathbb{R}^T \to \{0, \frac{1}{N}, \frac{2}{N}, \ldots, 1\}$. We will discuss the so-called smoothed Monte Carlo method, such that finding the optimal and feasible S given a timing vector Y becomes a Nonlinear Programming (NLP) problem.

4.1   MC-MILP optimization of the YS policy

The sample based approach for the YS policy can be handled by adding to the SP model a sample index $r = 1, .., N$ to the variables, I_{jtr} and Q_{tr} such that one has

replicas of the same variables that describe the actions of the model under each sample r. Furthermore, for the chance constraints one adds a binary variable $\delta_{tr} \in \{0,1\}$ representing the indicator value that specifies whether demand is fulfilled in period t in sample r

$$-I_{1tr} \leq m_t(1 - \delta_{tr}) \quad r = 1, \ldots, N, \; t = 1, \ldots, T \tag{12}$$

where m_t is an upper bound on the value of the out of stock $-I_{1t}$. This defines a function $\hat{a}_t(S) : \mathbb{R}^n \to \{0, \frac{1}{N}, \frac{2}{N}, \ldots, 1\}$ representing the reached service level under the set of samples. The corresponding chance constraints read

$$\hat{a}_t(S) := \frac{1}{N} \sum_{r=1}^{N} \delta_{tr} \geq \alpha, \; t = 1, \ldots, T. \tag{13}$$

The objective (1) is extended towards

$$\min \frac{1}{N} \sum_{t=1}^{T} \left(kY_t + \sum_{r=1}^{N} (cQ_{tr} + h \sum_{j=1}^{J-1} I_{jtr}^+ + wI_{Jtr}) \right), \tag{14}$$

with order quantity

$$Q_{tr} = (S_t - \sum_{j=1}^{J-1} I_{j,t-1,r})^+, \quad r = 1, \ldots, N, \; t = 1, \ldots, T \tag{15}$$

and the conventional order relation

$$S_t \leq \mathcal{M}Y_t \quad t = 1, \ldots, T \tag{16}$$

with a big-$\mathcal{M}$ value. The constraints (4) and (5) are extended to each sample

$$I_{1tr} = Q_{tr} - (d_{tr} - \sum_{j=1}^{J-1} I_{j,t-1,r})^+, r = 1, \ldots, N, t = 1, \ldots, T \tag{17}$$

and

$$I_{jtr} = (I_{j-1,t-1,r} - (d_{tr} - \sum_{i=j}^{J-1} I_{i,t-1,r})^+)^+, r = 1, \ldots, N, t = 1, \ldots, T, j = 2, \ldots, J. \tag{18}$$

Notice that the values that we intend to find, i.e. Y_t and S_t are independent of the sample r and the other variables that describe the simulation or evaluation part $Q_{rt}, I_{j,t,r}$ depend on the sample.

Solving the MC-MILP model is in most cases practically impossible due to the large number of binary variables δ and many solutions δ that represent the same obtained service levels $a(S)$. The number of samples $N = 5000$ mentioned in Section 3.3, implies defining for each period $N = 5000$ binary variables δ_{rt}. Instead, we will investigate a smoothed Monte Carlo approach as suggested in [3] to estimate the service levels in the MC-MILP model.

4.2   MC smoothing approach to the YS policy

First of all, consider the MC-MILP problem from the point of view of a NLP problem in the continuous variables S when order timing Y is given. The function $\hat{a}_t(S) : \mathbb{R}^n \to \{0, \frac{1}{N}, \frac{2}{N}, \ldots, 1\}$ in (13) and objective (14) are evaluated by using N sample paths following the dynamics (16), (17) and (18). The difficulty of applying an NLP solver for this problem is that (13) is piecewise constant, i.e. changing the values of S a bit may not change the evaluated value of $\hat{a}_t(S)$.
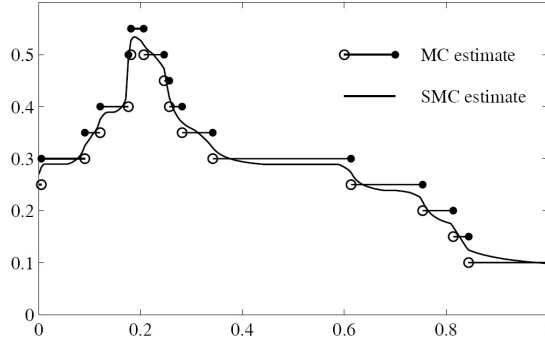
Fig. 1. Illustration of SMC from [3], where the estimated probability on the y-axis depends on varying one parameter on the x-axis

[3] show that one can make the reached service level practically a continuous function by following the MC smoothing approach. Let $z_{rt} = \sum_{j=1}^{J} I_{jtr}$ represent the total amount of product left over at the end of period t in sample r. One can measure, how close $\hat{a}_t(S)$ is to change value by the value of the least nonnegative total inventory $p_t^{[in]}(S) = \min_r\{z_{rt} | z_{rt} \geq 0\}$ and the least negative inventory $p_t^{[out]}(S) = \min_r\{-z_{rt} | z_{rt} < 0\}$. The suggested smoothing function $o_t(S)$ is

$$o_t(S) = \frac{1}{2N}\left(\frac{2p_t^{[in]}(S)}{p_t^{[in]}(S) + p_t^{[out]}(S)} - 1\right). \tag{19}$$

It is proven in [3], that $\hat{a}_t(S) + o_t(S)$ is continuous in the interesting values of S, as illustrated in Figure 1. Moreover, the function $\hat{a}_t(S) + o_t(S)$ deviates at most $\frac{1}{2N}$ from the reached service level $\hat{a}_t(S)$. This deviation is much smaller than the possible estimation error. Using $\hat{a}_t(S) + o_t(S)$ defines the problem $NLPS(Y)$ where constraint (13) in MC-MILP is replaced by

$$\hat{a}_t(S) + o_t(S) \geq \alpha, \ t = 1, \ldots, T \tag{20}$$

as a smooth optimization problem that in principle can be solved by a nonlinear optimization routine. Notice again that only values S_t have to be determined for $Y_t = 1$ and $\exists i = 1, \ldots, J-1, Y_{t-i} = 1$. For the chance constraints, one

only has to focus on the last period of the replenishment cycle $t + R_t - 1$; the demand in between will have a higher probability to be fulfilled. As starting point for the variables S_t in the nonlinear optimization the values $\hat{S}_{R_t,t}$ can be used. Algorithm 1 provides a list of order timing Y^* and order-up-to levels S^*

Algorithm 1 YSsmooth in: samples d_{tr}, cost data, $\alpha, \hat{S}_{R,t}$, out: Y^*, S^*

Set the best function value $C^U := \infty$
Generate a set of feasible order timing Y
for all Y
    if for the lower bound on cost $LBc(Y) < C^U$
        solve $NLPS(Y)$ using $\hat{S}_{R,t}$ values $\rightarrow S$ and cost C
        if $C < C^U$
            save the best found values $C^U := C, S^* := S, Y^* = Y$

---

that fulfils the chance constraints arbitrarily close if the number of samples N increases. One can use a lower bound on cost to decide that Y cannot be optimal. A lower bound $LBc(Y)$ on the cost contains the necessary minimum procurement cost $k \sum Y_t + c \sum E(d_t)$. Moreover, the expected inventory at the beginning of a period where no order takes place is at least $\hat{S}_{1,t}$ and the corresponding inventory cost can be added to the lower bound $LBc(Y)$. In an enumeration of Y, if $LBc(Y)$ is greater than the best feasible objective value C^U found so far, Y cannot be the optimal timing.

5   YQ(X): timing Y, stock-age dependent $Q_t(X)$

The YQ(X) policy fits nearly directly to the results found in Section 3. The decision maker is provided with an order timing vector Y. For each order moment $(Y_t = 1)$, the suggested order quantity Q_t depends on the age distribution X of the items in stock. For the order moments fulfilling the conditions of Lemmas 2 and 3, the order quantities are provided by using the pre-calculated basic order-up-to levels $\hat{S}_{R,t}$. For the other order moments, Lemma 4 tells us what to do when confronted with a negative inventory X. For a positive inventory X, the sample based estimation of Section 3 can be used. At an order moment, i.e. $Y_t = 1$ where the inventory position is positive and $R_t > 1$, the order quantity may be larger than the basic order-up-to level

$$Q_t(X) \geq (\hat{S}_{R_t,t} - \sum_{j=1}^{J-1} X_j)^+ \tag{21}$$

due to the expected out dating of inventory during the replenishment cycle. In order to determine the optimal order quantity for this case, we have to investigate the total inventory at the end of the replenishment cycle as function of the starting inventory X, the order quantity Q and the demand $d_t, \ldots, d_{t+R-1}$ during the replenishment cycle.

Definition 2. *The function $Z : \mathbb{R} \times \mathbb{R}^J \times \mathbb{R}^R \to \mathbb{R}$ is defined as the transformation $z = Z(Q, X, d)$ giving the total inventory $z = \sum_{j=1}^{J} I_{j,R-1}$ following the dynamics (4), (5) with starting inventory X, order quantity Q and demand vector $(d_1, \ldots, d_{R-1})$.*

This definition facilitates writing the order quantity we are looking for as the minimum value Q_t for which the chance constraint holds; this is the value of Q_t for which $P(Z(Q_t, X_t, \boldsymbol{d}_t, \ldots, \boldsymbol{d}_{t+R-1}) \geq 0) = \alpha$. The following property of function Z is useful.

Lemma 5. *Let function Z be defined by Definition 2, $R \leq J$, values for Q, X, d given. Let $Z(Q, X, d) = z$, then $\forall q \in \mathbb{R}$, $Z(Q + q, X, d) = z + q$.*

Proof. Due to $R \leq J$ following the equations (4), (5), non of the order quantity Q will be wasted (outdated). An additional amount q will be added to the total end inventory z. $\qquad \square$

Theorem 1. *Let function Z be defined by Definition 2, $R \leq J$, starting inventory X given, $\boldsymbol{z} = Z(0, X, \boldsymbol{d}_t, \ldots, \boldsymbol{d}_{t+R-1})$ with cdf Γ. The optimal order quantity in period t is $Q_t = (-\Gamma^{-1}(1 - \alpha))^+$.*

Proof. If $\Gamma^{-1}(1 - \alpha) > 0$, the current stock is enough to fulfil demand with α probability: $P(\boldsymbol{z} \leq 0) < (1 - \alpha) \to P(\boldsymbol{z} \geq 0) > \alpha$. So in that case, $Q_t = 0$ is optimal. For a value $Q_t = -\Gamma^{-1}(1 - \alpha) \geq 0$, we have $P(Z(0, X, \boldsymbol{d}_t, \ldots, \boldsymbol{d}_{t+R-1}) \leq -Q_t) = 1 - \alpha$. This implies $P(Z(0, X, \boldsymbol{d}_t, \ldots, \boldsymbol{d}_{t+R-1}) + Q_t \geq 0) = \alpha$. Using Lemma 5, this translates to $P(Z(Q_t, X, \boldsymbol{d}_t, \ldots, \boldsymbol{d}_{t+R-1}) \geq 0) = \alpha$. So the order quantity Q_t is the minimum value for which the end inventory has a probability of α to be positive. Therefore it is the optimal value. $\qquad \square$

Algorithm 2 $\mathbf{YQ}$(in: d_{rt}, cost data, α, $\hat{S}_{R,t}$), out: Y^*

$C^U := \infty$
Generate a set of feasible order timing Y
for all Y
  if for the lower bound on cost $LBc(Y) < C^U$
     Determine C by simulating N sample paths
     During the simulation
     if $Y_t = 1$
        if starting inventory X not positive or $R_t = 1$ take $Q_t = \hat{S}_{R_t,t} - \sum_{j=1}^{J-1} X_j$
        else simulate the replenishment cycle with N paths from X
        Determine the order quantity Q_t from (23)
  if $C < C^U$
     $C^U := C, Y^* = Y$

Lemmas 2 and 3 discussed the cases where $Q_t = \hat{S}_{R_t,t} - \sum_{j=1}^{J-1} X_j$ is the optimal solution. A possible deviation from this value in other cases is due to the waste

that can occur during the replenishment cycle. Taking this value as benchmark provides a corollary which follows directly from Theorem 1 and Lemma 5.

Corollary 1. *Let function Z be defined by Definition 2, $R \leq J$, starting inventory X given, $\boldsymbol{z} = Z(\hat{S}_{R_t,t} - \sum_{j=1}^{J-1} X_j, X, \boldsymbol{d}_t, \ldots, \boldsymbol{d}_{t+R-1})$ with cdf Γ. The optimal order quantity is $Q_t = \left(\hat{S}_{R_t,t} - \sum_{j=1}^{J-1} X_j - \Gamma^{-1}(1-\alpha)\right)^+$.*

In other words, Lemmas 2 and 3 discuss cases where the choice $Q_t = \hat{S}_{R_t,t} - \sum_{j=1}^{J-1} X_j$ gives $\Gamma^{-1}(1-\alpha) = 0$. Notice, this is also the case if the starting inventory is non-positive, $X_1 \leq 0$, as no waste can be generated. In other cases, waste can be generated and $\Gamma^{-1}(1-\alpha) < 0$. No analytical form is available to evaluate its value. To estimate the quantile $\Gamma^{-1}(1-\alpha)$, Monte Carlo simulation can be used as discussed in Section 3.3. Let D be an $N \times T$ matrix with samples $d_{r,t}$. For a starting inventory X, giving the order quantity $Q = \hat{S}_{R_t,t} - \sum_{j=1}^{J-1} X_j$, one can evaluate $z_r = Z(Q, X, d_{r,t}, \ldots, d_{r,t+R-1})$ being the total inventory of sample r at the end of the cycle. The adjusting amount $-\Gamma^{-1}(1-\alpha)$ is estimated by

$$A_t(X) = (-\text{quantile}(\{z_r, r = 1, \ldots, N\}, 1-\alpha))^+, \tag{22}$$

where quantile($\{\}, \alpha$) is the α sample quantile of set $\{\}$.

The order quantity for any starting inventory according to Corollary 1, can be approximated by

$$Q_t(X) = \hat{S}_{R_t,t} - \sum_{j=1}^{J-1} X_j + A_t(X). \tag{23}$$

This way of approaching the chance constraint is slightly stricter than the original service level constraints. It forces an α probability on positive inventory from any starting inventory X. This is also called a conditional service level constraint, see [8].

The order quantities for the YQ(X) policy are now defined either by the theoretical results, or by the sample based estimation in (22) and (23). The next question is to generate the best advice for the order timing Y. Algorithm (2) enumerates the possible timing vectors. Here one can make use of Lemma 1 and leave out those with too large periods between two orders. For each vector Y, the average cost is evaluated for a large simulation run that uses different random numbers than the ones in matrix D that are used to determine the order quantities by (22) and (23).

The YQ(X) policy has the advantage that it takes the age distribution into account. However, for the decision maker the required use of tables and possibly interpolation is more hassle than using a simple order-up-to strategy with a list of order-up-to levels of the YS policy. One question is in which cases the policy YQ(X) performs significantly better than the YS policy. The theoretical results already showed that in fact the YQ(X) policy works with basic order-up-to levels $\hat{S}_{R,t}$ for many cases. If costs and demand data are such that we order each period, or alternatively order every J periods, then in fact the YQ(X) policy works with order-up-to levels all the time according to Lemmas 2 and 3.

6   Numerical elaboration

Not taking the age distribution into account (YS policy) provides an easy to interpret policy. The developed method of Algorithm 1 aims to generate better solutions for the YS policy than the MILP approximation of [6] . Taking the age distribution into account (YQ(X) policy) should even lead to smaller expected costs. We used the experimental design of [6] to investigate the quality of the described policies in terms of how well the required service level is met and what are the expected costs. For which instances does it pay the trouble to take the age distribution into account?

Table 1. Base case with expected demand μ_t. Average cost and reached service level $\hat{\alpha}$ measured by simulating with 5000 runs

t\Cost	demand	MILP [6] 28882		YS Policy 28649		YQ(X) Policy 28205	
	μ_t	S_t	$\hat{\alpha}$	S_t	$\hat{\alpha}$	Y	$\hat{\alpha}$
1	800	1129	94.7%	1129	94.7%	1	1
2	950	1550	99.5%	1550	99.5%	0	98.7%
3	200	0	95.4%	0	95.4%	0	95.2%
4	900	2350	1	2340	1	1	1
5	800	0	98.7%	0	98.5%	0	98.7%
6	150	0	95.3%	0	94.7%	0	95.3%
7	650	1874	1	1874	1	1	1
8	800	0	95.3%.	0	95.3%	0	96.1%
9	900	1271	95.2%	1278	95.2%	1	94.9%
10	300	1333	1	1426	1	1	1
11	150	0	1	0	1	0	1
12	600	0	88.5%	0	95.1%	0	95.1%

We start with the base case with $k = 1500, c = 2, h = .5$, $w = 0$ $\alpha = 95\%$ from [6] and compare the implications of the reported S_t values by the MILP approach in that paper. The expected demand μ_t given in Table 1 for a time horizon of $T = 12$, and its variance is given by $(cv \times \mu_t)^2$ with variation coefficient $cv = 0.25$. The starting inventory is zero.

Algorithm 1 is used to generate the order-up-to levels S_t of the YS policy. Algorithm 2 generates the timing for the YQ(X) policy. With respect to efficiency, both algorithms required the order of magnitude of 5 minutes in a Matlab implementation. Taking into account the zero starting inventory (the first period we have to order), the number F_T of feasible timings Y in this case is 927. Both algorithms also remove more than 200 vectors due to cost bounding.

It is interesting to see that the YS policy (also elaborated in the MILP model) starts by choosing $Q_1 = S_1 = \hat{S}_{1,1}$ and the YQ(X) policy chooses $Q_1 = \hat{S}_{3,1}$, so aiming at covering demand of three periods. Although these numbers are exact, the numerical estimation of the service level via the simulation gives an error within the accuracy as $\hat{\alpha}_1 = 94.7\%$ instead of the used $\alpha = 95\%$. For this specific

case, notice that the YS and YQ(X) policies differ in number of orders in the time horizon and both provide a fulfilment of the service level constraints up to the simulation accuracy. This is in contrast to the reported result of the MILP model in [6], where for the last period the α probability was not reached. The YQ(X) policy is 2.4% cheaper than the YS policy for this case.

The illustrative case has been designed such that there is a large difference in cost between the YS and YQ(X) policy. To investigate the question when a large difference occurs, we repeated all 81 experiments of [6] and found the following. With regard to the disadvantage of using the more practical YS policy instead of taking the age distribution into account

- No disadvantage if cost structure implies ordering every period, or for a cycle that has the length of the shelf life.
- The disadvantage does not necessarily grow with the forecast error, i.e. larger uncertainty in demand.
- The disadvantage is less when the demand pattern is more stationary.

7   Conclusions

We investigated how order policies can be generated for a specific chance constrained inventory model where the order times should be fixed in advance for a finite horizon. Focus is on the idea that the chance constraints of the model can be approximated by an approach using Monte Carlo (MC) sampling.

Several theoretical properties have been derived for the optimal order quantities. For those order moments where the analytical results do not apply, sample based estimation of the service level can be used. For the YQ(X) policy, the order quantity is minimized such that from each starting inventory for a simulated replenishment cycle the chance constraint is just fulfilled. We derived an algorithm to determine the optimal order timing for this order policy.

When investigating the possibility of generating order-up-to levels for a so-called YS policy, one can in principle formulate a MC-MILP problem that is usually unsolvable. We show that such a problem can be approximated arbitrarily near by using the smoothed Monte Carlo method to construct a MINLP problem that is integer in the timing variables Y and continuous in the order-up-to variables S. For each timing vector Y to be evaluated, a continuous $NLPS(Y)$ problem should be solved. A specific algorithm based on enumeration and bounding has been derived to solve the problem.

The YQ(X) policy is more complicated for providing decision support to the decision maker. The derived algorithms and theoretical results have been used to investigate the question for which instances the easier YS policy provides (nearly) the same performance as the more complicated YQ(X) policy. Mainly a wide variation in the demand pattern over the planning horizon provides an advantage for the YQ(X) policy that takes the age distribution into account in the determination of the order quantity.

References

1. A.G. Alcoba, E.M.T. Hendrix, I. García, G. Ortega, K.G.J. Pauls-Worm, and R. Haijema. On computing order quantities for perishable inventory control with non-stationary demand. In B. Murgante et al., editor, *Proceedings of ICCSA 2015*, Lecture Notes in Computer Science, Cham, 2015. Springer.
2. J.R. Birge and F. Louveaux. *Introduction to Stochastic Programming.* Springer, New York, 1997.
3. E. M. T. Hendrix and N. J. Olieman. The smoothed Monte Carlo method in robustness optimisation. *Optimization Methods and Software*, 23:717–729, 2008.
4. Jean Jacod and Philip Protter. *Probability Essentials (2ed).* Springer-Verlag, Berlin, 2004.
5. R. Lyman Ott and Michael Longnecker. *An Introduction to Statistical Methods and Data Analysis (5th ed.).* Duxburry, Pacific Grove, CA 93950 USA, 2001.
6. K. G. J. Pauls-Worm, E. M. T. Hendrix, R. Haijema, and J. G. A. J. van der Vorst. An milp approximation for ordering perishable products with non-stationary demand and service level constraints. *International Journal of Production Economics*, 157:133–146, 2014.
7. W.A. Rijpkema, E.M.T. Hendrix, R. Rossi, and J.G.A.J. van der Vorst. Application of stochastic programming to reduce uncertainty in quality-based supply planning of slaughterhouses. *Annals of Operations Research*, to appear, 2013.
8. R. Rossi, S.A. Tarim, B. Hnich, and S. Prestwich. A global chance-constraint for stochastic inventory systems under service level constraints. *Operations Research*, 13:490–517, 2008.
9. E. A. Silver, D. F. Pyke, and R. Peterson. *Inventory Management and Production Planning and Scheduling.* Wiley, 1998.
10. Eylem Tekin, Ülkü Gürler, and Emre Berk. Age-based vs. stock level control policies for perishable inventory systems. *European Journal of Operational Research*, 124:309 – 329, 2001.