WORKING PAPER
Wageningen School of Social Sciences

On a Global Optimization solution approach for a stochastic programming inventory control model

Eligius M.T. Hendrix, Karin G.J. Pauls-Worm, Rene Haijema, Roberto Rossi

WASS Working PAPER No. 8
2013

Wageningen School
of Social Sciences

Hollandseweg 1, 6706 KN Wageningen,
The Netherlands
Phone: +31 317 48 41 26
Fax: +31 317 48 47 63
Internet: http://www.wass.wur.nl
e-mail: wass@wur.nl

Eligius M.T. Hendrix	Wageningen University, eligius.hendrix@wur.nl
Karin G.J. Pauls-Worm	Wageningen University, karin.pauls@wur.nl
Rene Haijema	Wageningen University, rene.haijema@wur.nl
Roberto Rossi	University of Edinburgh, robros@gmail.com

On a Global Optimization solution approach for a stochastic programming inventory control model

Eligius M.T. Hendrix, Karin G.J. Pauls-Worm, Rene Haijema
Wageningen University, eligius.hendrix@wur.nl, karin.pauls@wur.nl,rene.haijema@wur.nl

Roberto Rossi
University of Edinburgh, robros@gmail.com

The basis of this study is a Stochastic Programming (SP) model derived for a practical case of a specific inventory control problem for a perishable product. As it contains chance constraints describing the service level, deriving policies for this model is a challenge. In order to find parameter values in an order-up-to level policy, we derive a conventional Monte Carlo based MC-MILP model. It is shown that this problem is practically equivalent to a nonconvex MINLP model with binary variables describing when to order and continuous variables providing the order-up-to levels. A specific algorithm is designed to solve the problem by enumeration and bounding and iterative nonconvex nonlinear optimization. A second policy is outlined that takes the age distribution in the model into account. After deriving the optimal timing of the orders, every order quantity is generated by sampling the review period. As the MINLP based policy is easier for the decision maker, our question is for which cases that policy is sufficient.

Key words: Inventory control, Perishable products, MINLP, stochastic programming, Monte Carlo

1. Introduction

The basis of our study is a SP model published in Pauls-Worm et al. (2013) 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 demand is uncertain and non-stationary such that one produces to stock. The 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 $t \in \{1, 2, ..., T\}$. The latter implies being mathematically confronted with a chance constraint. It is known, that this may lead to challenging Global Optimization problems. For a recent overview, see Parpas et al. (2009) and the references therein.

The solution for such a model is a so-called order policy. Given the inventory situation $\mathbf{I}$ at the beginning of period moment t, an order policy should advice the decision maker on the order quantity Q_t. For the decision maker, simple rules are preferred. Therefore, Pauls-Worm et al. (2013) generate a simple so-called YS policy that consists of a list of order periods Y with order-up-to levels S_t, i.e. each order period the manager replenishes the inventory to a level S_t. An MILP approximation is described in that paper that provides a list S_t in a short time, but unfortunately it does not exactly fulfill the service level constraints for all instances. Therefore, the 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. We show that finding the timing Y and such values for S based on (smoothed) Monte Carlo sampling requires an MINLP problem to be solved. 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.

Moreover, we investigate the generation of a policy that also looks for the best order moments Y to fix, but in the decision of the order quantity Q_t takes the distribution of the age of items in stock into account. The decision support enhances a far more complicated rule, as now a simple list of levels is not sufficient anymore. Our question is how to generate such a policy and whether the simple policy is doing much worse than this policy.

The effectiveness of a policy can be investigated by simulation. Using pseudo-random series of the demand, one measures how well the required service levels are met and estimate the expected cost. This paper is organised as follows. Section 2 describes the underlying SP model with chance constraints. Section 3 describes the generation of the quantities S_t via the MC-MILP and MINLP problems and Section 4 provides the elaboration of a policy that does take the age distribution into account. In Section 5, we numerically compare the performance of the policies derived from the different approaches with several instances. Section 6 summarizes our findings.

2. Stochastic Programming Model

The stochastic demand implies that the model has random inventory variables 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

$$\mathbf{I} = (I_{1,t-1}, \ldots, I_{J-1,t-1}), \tag{1}$$

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 $x^+ = max\{x, 0\}$. The formal SP model is given in detail.

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
- d_t Normally distributed demand with expectation μ_t and variance $(cv \times \mu_t)^2$ where cv is a given coefficient of variation.
- k fixed ordering cost
- c procurement cost
- h inventory cost
- w disposal cost, is negative when having salvage value
- α service level

Variables
- $Q_t \geq 0$ ordered and delivered quantity at beginning period t
- I_{jt} Inventory of age j at end of period t, initial inventory fixed $I_{j0} = 0$, I_{1t} free, $I_{jt} \geq 0$ for $j = 2, \ldots, J$.

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

$$E\left(\sum_{t=1}^{T}\left(h\sum_{j=1}^{J-1} I_{jt}^{+} + g(Q_t) + wI_{Jt}\right)\right) = \sum_{t=1}^{T} E\left(g(Q_t) + h\sum_{j=1}^{J-1} I_{jt}^{+} + wI_{Jt}\right), \tag{2}$$

where procurement cost is given by the function

$$g(x) = k + cx, \quad \text{if} \quad x > 0, \text{ and } \quad g(0) = 0. \tag{3}$$

The chance constraint expressing the required service level is

$$P(I_{1t} \geq 0) \geq \alpha, \ t = 1, \ldots, T \tag{4}$$

and the dynamics of the inventory of the items of different ages is described by

$$I_{1t} = Q_t - (d_t - \sum_{j=1}^{J-1} I_{j,t-1})^{+}, \ t = 1, \ldots, T \tag{5}$$

and

$$I_{jt} = \left(I_{j-1,t-1} - (d_t - \sum_{i=j}^{J-1} I_{i,t-1})^+ \right)^+ , \; t = 1, \ldots, T, j = 2, \ldots, J. \tag{6}$$

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. A feasible order policy $Q_t(\mathbf{I})$ of the SP model fulfils the nonnegativity aspects and equations (4), (5) and (6). An optimal policy also minimises (2).

For a nonperishable, a common way to deal with the planning is to define so-called order-up-to levels S_t, e.g. Silver et al. (1998). Each period where a replenishment takes place, one aims to cover the demand of $R(t)$ periods. This policy is easy to grasp for the decision maker and is called here a policy with a replenishment cycle $R(t)$ and stock age independent order-up-to level S_t. 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 Section 3.

In a replenishment cycle concept, taking the age distribution into account, requires supplying the decision maker with an information system that advices on the order quantity $Q_t(\mathbf{I})$. In Section 4, we investigate how order quantities can be derived. As this policy is wider, it should provide lower expected cost than the case where the age distribution is not taken into account.

3. YS: replenishment cycle $R(t)$, stock age independent order-up-to level S_t

Literature on inventory control e.g. Silver et al. (1998) applies the concept of a replenishment cycle, i.e. the length of the period R for which the order of size Q is meant. For non-stationary demand, the replenishment cycle $R(t)$ depends on the period. We model this here as to provide the decision maker a list $Y \in \{0,1\}^T$ of periods when to order, such that in fact $Y_t = Y_{t+R(t)} = 1$ and $Y_k = 0, k = t+1, .., t+R(t)-1$. Another important concept is that of the so-called order-up-to level S_t. When ignoring the age distribution of the inventory, the decision maker replenishes in the order period the inventory up to a level S_t:

$$Q_t(\mathbf{I}) = (S_t - \sum_{j=1}^{J-1} \mathbf{I}_j)^+, \; t = 1, \ldots, T, \tag{7}$$

where values for the order-up-to level S_t should be determined. If no order takes place, one can define $S_t = 0$ which co-incides with $Y_t = 0$.

In the determination of the best values, we first fix the list of order periods Y and then try to find the best values for S_t. This seems an easy problem for a nonperishable product, i.e. no waste is generated. A replenishment cycle starting at period t of length $R(t)$ is dealing with a total demand $d_t + .. + d_{t+R(t)-1}$. Let $G_{t,R(t)}$ be the cumulative distribution function (cdf) of this total demand. The chance constraint is fulfilled by taking as order-up-to level the value $\sigma_{t,R} = G^{-1}_{t,R}(\alpha)$. One of the challenges is that even this level may be too high in a model with non-stationary demand. Namely, the order quantity in the former cycle may have been so high, that there is a positive probability that also the current replenishment period is covered. In that case, the order-up-to level can be taken a bit lower than $\sigma_{t,R}$.

Moreover, for a perishable product, one has to take into account that items in stock may perish and one may like to order more than that level, i.e. $S_t \geq \sigma_{t,R}$. Given Y, we are dealing with an NLP problem $NLP(Y)$ in the continuous variable S minimising (2) with constraints (4), (5), (6) and (7). The difficulty is, that the chance constraint (4) is not an analytical expression in S that can easily be evaluated. Therefore we will discuss in Sections 3.1 and 3.2, how Monte Carlo simulation may be used as an approximation.

The next question is how to deal with the integer part Y of the solution. Pauls-Worm et al. (2013) provide an approximate strategy based on a MILP model that not only takes the quantities $\sigma_{t,R}$ into account, but also the expected waste and age distribution. The elegance of the approach is that no dedicated software is used other than standard commercially available MILP solvers that generate in a short time a list Y of order moments and corresponding order-up-to levels S. The procedure provides service levels that approximate the requirement α, but for some instances the levels do not reach this requirement in some of the periods.

To find the optimal values fulfilling the constraints, one can in principle solve $NLP(Y)$ for all feasible Y and select the best one. Notice that $Y_1 = 1$ in the model, as we start with zero inventory leaving 2^{T-1} binary vectors representing different timing of the replenishment. The case we are interested in concerns $T = 12$ periods, so there are 2048 possible timing vectors for ordering. In the search for the optimal timing and for solving $NLP(Y)$, the following reasoning can be used.

• Review periods $R > J$ can not take place due to perishability. Series Y with J zeros in a row are not feasible. For $J = 3, T = 12$ this leaves 927 vectors Y.

• One can use a lower bound on cost to decide that Y cannot be optimal. As lower bound $LBc(Y)$ on the cost of $NLP(Y)$, one can take the necessary minimum procurement cost $k\sum Y_t + c\sum E(d_t)$ and a minimum inventory cost knowing that at the beginning of the period where no order takes place, at least $\sigma_{t,1}$ has to be available. 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.

• The variables S_t in $NLP(Y)$, only correspond to the periods where $Y_t = 1$. However, if no inventory is available, i.e. $\mathbf{I} = 0$, due to being the first period, or the former replenishment cycle has length J, then one orders the minimum necessary amount $S_t = \sigma_{t,R}$.

This leaves us with solving $NLP(Y)$ for many timing vectors Y, where constraints (4) are not an analytical expressions in S. How to deal with that?

3.1. MC-MILP approximation of the YS policy

A usual way to deal with chance constraints in stochastic programming is called scenario-based modeling, Birge and Louveaux (1997). This enhances to use Monte Carlo (MC) simulation and generate N samples d_{tr} of the demand series d with pseudo-random numbers. For the objective function this provides a simple extension, as it is expressed in the continuous variables only. However, it is known from literature Hendrix and Olieman (2008), that when using samples to measure a probability (service level) the resulting function is piecewise constant in the continuous variable, here S_t.

Conceptually, one can also define the problem in MILP terms adding in the SP model a simulation run index $r = 1,..,N$ to the variables, I_{jtr}, Q_{tr} and adding a binary variable δ_{tr} that specifies whether demand is fulfilled in period t in run r

$$-I_{1tr} \leq= m_t(1-\delta_{tr}) \quad r = 1,\dots,N,\ t = 1,\dots,T \tag{8}$$

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

$$a_t(S) := \frac{1}{N}\sum_k \delta_{tr} \geq \alpha,\ t = 1,\dots,T. \tag{9}$$

The objective (2) is extended towards

$$\frac{1}{N}\sum_{t=1}^{T}\sum_{r=1}^{N}\left(g(Q_{tr}) + h\sum_{j=1}^{J-1} I^+_{jtr} + wI_{Jtr}\right), \tag{10}$$

with order quantity

$$Q_{tr} = Y_t \times (S_t - \sum_{j=1}^{J-1} I_{j,t-1,r})^+, \quad r = 1,\dots,N,\ t = 1,\dots,T. \tag{11}$$

The constraints (5) and (6) are extended to each run

$$I_{1tr} = Q_{tr} - (d_{tr} - \sum_{j=1}^{J-1} I_{j,t-1,r})^+, \; r = 1,\dots,N, \; t = 1,\dots,T \tag{12}$$

and

$$I_{jtr} = \left(I_{j-1,t-1,r} - (d_{tr} - \sum_{i=j}^{J-1} I_{i,t-1,r})^+ \right)^+, r = 1,\dots,N, t = 1,\dots,T, j = 2,\dots,J. \tag{13}$$

Rewriting (11) and the function $()^+$, the complete model can be expressed in MILP terms. However, solving is practically impossible due to the large number of binary variables δ and many solutions δ that represent the same obtained service levels $a(S)$. Instead, we will investigate a Monte Carlo smoothing approach as suggested in Hendrix and Olieman (2008) following the MC-MILP model.

3.2. MC smoothing approach to the YS policy

We focus again on $NLP(Y)$ where the function $a(S)$ in (9) and objective (10) are evaluated by a simulation of N sample paths following the dynamics (11), (12) and (13). Let us first remark that the constraints (9) only have to be evaluated for the critical periods at the end of a replenishment cycle, reducing their number to $\sum Y_t$. The difficulty of applying an NLP solver for this problem is that (9) is piecewise constant, i.e. changing the values of S a bit may not change the evaluated value of $a_t(S)$. However, one can make the approximate reached service level practically a continuous

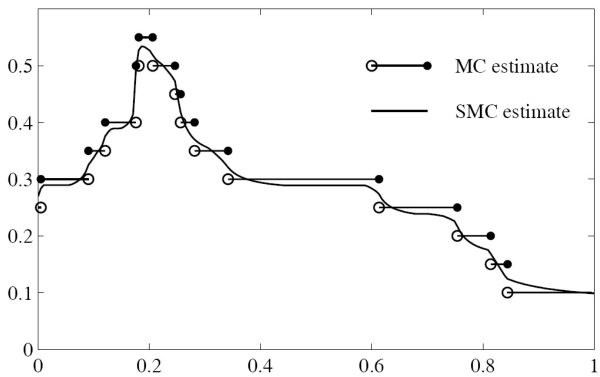

Figure 1 Illustration of SMC from Hendrix and Olieman (2008), where the estimated probability on the y-axis depends on varying one parameter on the x-axis

function by following the MC smoothing approach proposed in Hendrix and Olieman (2008). Let $Itot_{tr} = \sum_{j=1}^{J} I_{jtr}$ represent the total amount of product left over at the end of period t in run r. One can measure, how close $a(S)$ is to change value by the value of the least nonnegative total inventory $p_t^{[in]}(S) = \min_r\{Itot_{tr} | Itot_{tr} \geq 0\}$ and the least negative inventory $p_t^{[out]}(S) = \min_r\{-Itot_{tr} | Itot_{tr} < 0\}$. The suggested smoothing function $s_t(S)$ is

$$s_t(S) = \frac{1}{2N} \left(\frac{2p_t^{[in]}(S)}{p_t^{[in]}(S) + p_t^{[out]}(S)} - 1 \right). \tag{14}$$

It is shown in Hendrix and Olieman (2008), that $a_t(S) + s_t(S)$ is continuous in the interesting values of S, as illustrated in Figure 2. This defines the problem $NLPS(Y)$ where constraint (9) in MC-MILP is replaced by

$$a_t(S) + s_t(S) \geq \alpha, \; t = 1,\dots,T \tag{15}$$

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 $R(t)+J-1$. As starting point for the variables S_t in the nonlinear optimization we use the values $\sigma_{t,R(t)}$. Algorithm 1 provides a list of order timing Y^* and order up to levels S^* that fulfils the

Algorithm 1 YSsmooth(in: samples d_{tr},cost data, α, σ_{tR}), out: $Y^*, S^*(x)$)

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 σ_{tR} values $\rightarrow S$ and cost C
    if $C < C^U$
      save the best found values $C^U := C, S^* := S, Y^* = Y$

chance constraints arbitrarily close if the number of samples N increases.

4. YQ: replenishment cycle $R(t)$ and stock-age dependent order level $Q_t(I)$

This policy is more general than the YS policy where now the replenishment moments Y are again fixed, but the suggested order quantity Q_t depends on the age distribution $\mathbf{I}$ of the items in stock. From a decision support viewpoint, now the decision maker requires more information than a simple order-up-to value list. $Q_t(\mathbf{I})$ is in fact a function that either has to be provided by appropriate software, or represented by a table.

Considering the properties of this order quantity function, notice that in case shortage appears, i.e. $\mathbf{I}_1 < 0$, the optimal quantity is $Q_t(\mathbf{I}) = Q_t(0) - \mathbf{I}_1 = \sigma_{tR(t)} - \mathbf{I}_1$. For the other values of the inventory, the values of Q_t in the table can be interpolated to obtain an advice on the appropriate order level $Q_t(\mathbf{I})$. We can follow again the approach of enumerating possible timing vectors Y. However, how to determine the order quantities Q_t in an order period, i.e. $Y_t = 1$ if the inventory position is positive

Algorithm 2 YQ(in: samples d_{tr},cost data, α, σ_{tR}), 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 $\mathbf{I}$ is not positive or $R(t) = 1$ take $Q_t = \sigma_{t,R(t)} - \sum_{j=1}^{J-1} \mathbf{I}_j$
      else simulate the replenishement period with N paths from $\mathbf{I}$
      Determine the order quantity Q_t from (18)
  if $C < C^U$
    $C^U := C, Y^* = Y$

inventory position and $R(t) > 1$? The order quantity may be larger than the safety level, i.e.

$$Q_t(\mathbf{I}) \geq L_t := (\sigma_{tR(t)} - \sum_{j=1}^{J-1} \mathbf{I}_j)^+. \tag{16}$$

due to the expected outdating of inventory during the replenishment cycle. One can estimate the expected waste during the replenishment cycle by MC simulation of the final inventory in a replenishment cycle and increase the order-up-to level with that in order to guarantee the service level. This works as follows.

Let $I_{t+R-1,1,r}$ be the simulated freshest product inventory of run r at the end of the cycle, then

$$EW(Q,\mathbf{I}) = (-quantile(\{I_{t+R-1,1,r}, r = 1,\ldots,N\}, 1-\alpha))^+ \tag{17}$$

estimates the expected waste from starting inventory position $\mathbf{I}$ given an order quantity Q, where $quantile(\{\},\beta)$ is the β sample quantile of set $\{\}$. So, if the simulation predicts a negative expected $(1-\alpha)$-quantile of the final inventory, all of it is due to the expected waste. We can now simply increase the order quantity with this amount to guarantee the service level:

$$Q_t(\mathbf{I}) = L_t + EW(L_t, \mathbf{I}). \tag{18}$$

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 $\mathbf{I}$. This is also called a conditional service level constraint, see Rossi et al. (2008). On the other hand, the possible order quantities are more free to be chosen than in the YS policy. Algorithm (2) shows how the MC procedure can be used to identify the best timing vector Y^*. The generated policy should in principle provide lower expected cost. The drawback from the decision maker point of view is that now he is provided a list of tables that require interpolation to derive a good order quantity. In the sequel our question is, how much can be saved by taking the age distribution into account.

If one also relaxes the requirement that the decision maker is provided with a list Y of order moments, we arrive at a policy where one can speculate on the probability of not having to order in the following period. In Hendrix et al. (2012) this policy is aimed for by applying a Stochastic Dynamic programming (SDP) approach.

5. Comparative Study

Not taking the age distribution into account (YS policy) provides an easier to interpret policy than taking the age distribution into account (YQ policy). The question is, what is 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? For the illustration, we start with a base case from Pauls-Worm et al. (2013) and also provide the resulting YS policy reported in that paper. In this case the costs are $k = 1500, c = 2, h = .5$ and $w = 0$. The required service level is $\alpha = 95\%$. The expected demand μ_t given in Table 1, and its variance is given by $(cv \times \mu_t)^2$ with variation coeficient $cv = 0.25$.

Algorithm 1 is used to generate the order-up-to levels S_t of policy 1. Algorithm 2 generates the timing for policy 2. With respect to efficiency, both algorithms required the order of magnitude of 5 minutes in a Matlab implementation. One should notice, that for Algorithm 1, $N = 1000$ samples are used, as this is sufficient to distinguish the best timing. In the simulation, for the choice of the order quantity $Q(\mathbf{I})$, simulated demand samples were based on 5000 numbers that differ from the series used in the MC simulation to evaluate the effectiveness. In both algorithms, more than 200 of the 927 feasible timings Y could be removed due to cost bounding.

Table 1 Base case with expected demand μ_t. Average cost and service level sl measured in a MC simulation, 5000 runs

Cost	demand	YS Policy 28882		MILP approx. 28649		YQ Policy 28205	
t	μ_t	S_t	sl	S_t	sl	Y	sl
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	2340	1	2350	1	1	1
5	800	0	98.5%	0	98.7%	0	98.7%
6	150	0	94.7%	0	95.3%	0	95.3%
7	650	1874	1	1874	1	1	1
8	800	0	95.3%.	0	95.3%	0	96.1%
9	900	1278	95.2%	1271	95.2%	1	94.9%
10	300	1426	1	1333	1	1	1
11	150	0	1	0	1	0	1
12	600	0	95.1%	0	88.5%	0	95.1%

For this specific case, notice that the policies differ in number of orders in the time horizon and both provide a sharp approximation of the service level constraints. This is in contrast to the reported result of the approximating MILP model in Pauls-Worm et al. (2013), where for the last period the α probability was not reached. The YQ policy is 2.4% cheaper than the YS policy for this case.

In which cases is there a large difference in cost between the YS and YQ policy? To investigate this question, we look into hypotheses and start varying the parameters of the base case one by one. No difference is expected when the order cost k are relatively high. Any policy will order as few as possible T/J times and no inventory goes from one replenishment cycle to the other; in our case one should order $Q_1 = \sigma_{1,3}, Q_4 = \sigma_{4,3}, Q_7 = \sigma_{7,3}, Q_{10} = \sigma_{10,3}$. For the base case, this policy appears for a order cost of $k = 2000$. Both algorithms generated this policy in about 2 minutes, where the cost bounding is relatively effective as $Y = (1,0,0,1,0,0,1,0,0,1,0,0)$ is the first feasible timing and ordering more time enhances higher order cost. The evaluation by simulation provides practically the same expected costs of 31300.

We vary the service level to $\alpha \in \{.9, .95, .98\}$. We expect that for the YQ policy which is more flexible, it is easier to reach higher service levels than for the YS policy. For $\alpha = .9$ both policies provide practically the same cost of 27900 whereas for $\alpha = .95$ the YQ policy was 2.4% cheaper. Increasing the service level to $\alpha = .98$ requires much larger safety stocks that let both policies choose for far more orders (7) providing in the end a cost difference of only 1% in favour of the YQ policy. From these experiements no strong trend follows that shows more advantage of taking the age distribution into account with increasing service level requirement.

To simulate the concept of having less or more uncertainty in the demand (forecast) distributions, we varied in the base case the coefficient of variation $cv \in \{.1, .25, .33\}$. Although we expected a similar behaviour as varying the required service level, as again the number of orders will increase with increasing uncertainty, we observed a trend for this case. For $cv = .33$ again the resulting YQ policy is 1% cheaper and both policies go for more orders. However, for less uncertainty with $cv = .1$ both policies again return to replenishment cycles of $J = 3$ periods and behave practically the same.

We observed that the cost advantage of 2.4% when taking the age distribution into account was the largest in the base case. As well for the case where every moment an order takes place as for the situation where replenishment periods are $R = J$ (high order cost, low uncertainty), one may as well use the YS policy. Does the advantage in the intermediate case hold if we have other demand patterns? To investigate this question, we use the three variants of the demand pattern in Pauls-Worm et al. (2013) that all have the same total expected demand of 7200 units. The patterns are depicted in Figure 1.

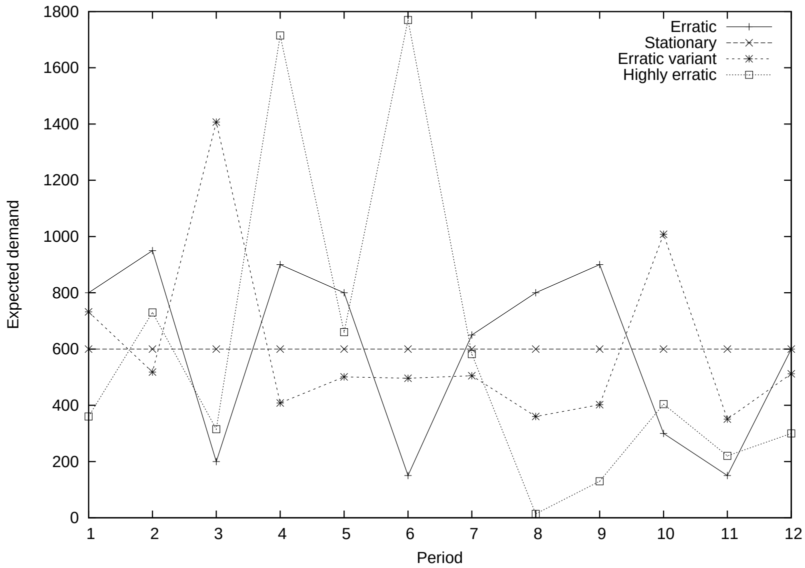

Figure 2 Different demand patterns with a total expected demand of 7200

For the stationary demand, both policies have replenishment cycles of two periods and provide the same rule; no difference is observed in the resulting cost. The variant of the erratic demand pattern of the base case was constructed to provide worst case behaviour of the MILP policy. The policies developed in this paper, however, despite differing in order moments, provide a good approximation of service levels and the YQ policy appears less than 1% cheaper. The highly erratic demand scenario attempts to model extreme case behaviour, where variation of expected demand between periods is more than 10 times the value of demand. For such an extreme case, indeed one observes a cost advantage of taking the age distribution into account. The advantage is 1000 over a cost of 28000, so more than 3.5%.

Concluding, instances where the cost structure implies ordering every period, or every J periods, one can as well apply the YS policy without any loss. Larger forecast errors, modelled here by the cv do not seem to favour the use of taking the age distribution into account. The advantage of applying the more difficult to comprehend YQ policy is more worth the trouble if the expected demand has a big variation over the periods.

6. Conclusions

We investigated how global optimization can be used to generate order policies 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 a Monte Carlo (MC) sampling approach.

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 practically unsolvable. We show that such a problem can be approximated arbitrarily near by using the smoothed Monte Carlo method to construct an MINLP problem that is integer in the timing variables Y and continuous in the order-up-to variables S. For each to evaluate timing vector, a nonconvex continuous $NLPS(Y)$ problem should be solved. A specific algorithm based on enumeration and bounding has been derived to solve the problem.

A second so-called YQ policy has been constructed where the order quantity is minimized such that from each starting inventory for a simulated replenishment cycle the chance constraint is just fulfilled. On one hand, this provides a more strict fulfillment of the chance constraint, but on the other hand, the order quantity is more free to be chosen than in the YS policy. This policy is harder to be provided to a decision maker.

Algorithms have been presented and implemented to generate the timing and order-up-to levels of the policies. The algorithms make use of the mathematical characteristics of the problem. The algorithms have been used to investigate the question for which instances the easier YS policy provides (nearly) the same performance as the more complicated YQ policy. Only for very extreme variation in demand over the planning horizon, the YQ policy that takes the age distribution into account saves more than 1% of the costs in the tested cases.

References

Birge, J.R., F. Louveaux. 1997. *Introduction to Stochastic Programming*. Springer, New York.

Hendrix, E. M. T., N. J. Olieman. 2008. The smoothed Monte Carlo method in robustness optimisation. *Optimization Methods and Software* **23** 717–729.

Hendrix, E.M.T., R. Haijema, R. Rossi, K. G. J. Pauls-Worm. 2012. On solving a stochastic programming model for perishable inventory control. B. Murgante, ed., *Proceedings of ICCSA 2012*, *Lecture Notes in Computer Science*, vol. 7335. Springer, Heidelberg, 47–65.

Parpas, P., B. Rustem, E.N. Pistikopoulos. 2009. Global optimization of robust chance constrained problems. *Journal of Global Optimization* **43**(2) 231–247.

Pauls-Worm, K. G. J., E. M. T. Hendrix, R. Haijema, J. G. A. J. van der Vorst. 2013. Inventory control for a perishable product. *International Journal of Production Economics* .

Rossi, R., S.A. Tarim, B. Hnich, S. Prestwich. 2008. A global chance-constraint for stochastic inventory systems under service level constraints. *Operations Research* **13** 490–517.

Silver, E. A., D. F. Pyke, R. Peterson. 1998. *Inventory Management and Production Planning and Scheduling*. Wiley.