On solving a stochastic programming model for perishable inventory control*

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

[1] Computer Architecture, Universidad de Málaga
[2] Operations Research and Logistics, Wageningen University
[3] Business School, University of Edinburgh

Abstract. This paper describes and analyses a Stochastic Programming (SP) model that is used for a specific inventory control problem for a perishable product. The decision maker is confronted with a non-stationary random demand for a fixed shelf life product and wants to make an ordering plan for a finite horizon that satisfies a service level constraint. In literature several approaches have been described to generate approximate solutions. The question dealt with here is whether exact approaches can be developed that generate solutions up to a guaranteed accuracy. Specifically, we look into the implications of a Stochastic Dynamic Programming (SDP) approach.
keywords: Stochastic Programming, Dynamic Programming, Inventory control, Perishable products, Service level constraint.

1 Introduction

We consider a 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. To keep waste due to outdating low, one issues the oldest product first, i.e. FIFO issuance. A service level applies to guarantee that the probability of not being out-of-stock is higher than α in every period $t \in \{1, 2, ..., T\}$. Any unmet demand is backlogged. In [9], a Stochastic Programming model has been introduced that uses a chance constraint to generate order policies for this planning problem. Specifically, an MILP re-formulation to generate an approximate solution and an enumeration method based on Sample Average Approximation are investigated in that paper.

The question is here, whether methods can be developed based on the properties of the model to generate solutions for a realistic time horizon and shelf life based on Stochastic Dynamic Programming (SDP). The target is to generate

* This paper has been supported by The Spanish Ministry of Science and Innovation (project TIN2008-01117) and Junta de Andalucía (P11-TIC-7176), in part financed by the European Regional Development Fund (ERDF). Eligius M.T. Hendrix is a fellow of the Spanish "Ramon y Cajal" contract program, co-financed by the European Social Fund.

solutions for the instances used in [9] which relates to a practical planning problem with a time horizon of $T = 12$ and shelf life of the perishable products of $J = 3$ periods. This is followed by questions on how solution approaches behave for varying setings of the problem.

This paper is organised as follows. Section 2 describes the underlying SP model. Section 3 describes a conceptual solution approach based on Stochastic Dynamic Programming (SDP) and the properties of the underlying problems to be solved. In Sect. 4 we illustrate the approach with several instances. Section 5 summarizes our findings.

2 Stochastic programming model

The model is summarised from an optimisation perspective. As much as possible, in the used symbols, we distinguish between model parameters (exogenous in lower case letters) and decision variables (capital letters) that are characterised by direct decision variables and dependent stock variables. Capitals are also used for upper bounds on the indices.

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 cdf Γ_t
k fixed ordering cost, euro
c procurement cost, euro/ton
h inventory cost, euro/ton
w disposal cost, euro/ton, is negative when having salvage value
α service level

Here we have assumed the product demand to be continuous, but the model can be modified easily to deal with the discrete demand case.

Variables

$Q_t \geq 0$ ordered and delivered quantity at beginning period t, ton
I_{jt} Inventory of age j at end of period t,
with the initial closing inventory levels fixed to $I_{j0} = 0$, and variable I_{1t} is free (as we assume backlogging and FIFO issuing), whereas $I_{jt} \geq 0$ for $j = 2, \ldots, J$.

The inventory variables (apart from the initial fixed levels I_{j0}) are random variables due to the stochastic nature of demand. If the order decision Q_t depends on the inventory levels $(I_{1,t-1}, \ldots, I_{J-1,t-1})$, Q_t is also a random variable. In the notations below $P(.)$ denotes a probability and $E(.)$ is the expected value operator. The total expected costs over the finite horizon is to be minimised:

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

where procurement cost is given by the function

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

The chance constraint is

$$P(I_{1t} \geq 0) \geq \alpha, \ \ t = 1, \ldots, T \tag{3}$$

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

$$I_{1t} = Q_t - (d_t - \sum_{j=1}^{J-1} I_{j,t-1})^+, \ \ t = 1, \ldots, T \tag{4}$$

and

$$I_{jt} = \left(I_{j-1,t-1} - (d_t - \sum_{i=j}^{J-1} I_{j,t-1})^+ \right)^+, \ \ t = 1, \ldots, T, j = 2, \ldots, J. \tag{5}$$

These dynamics equations describe the FIFO issuing policy and imply that I_{1t} is a free variable, whereas I_{jt} is nonnegative for the older vintages $j = 2, \ldots, J$. A feasible order strategy $Q_t(I)$ of the SP model fulfills the nonnegativity aspects and equations (3), (4) and (5). An optimal solution also minimises (1).

The approaches described in [9], in fact look for the best static moments of ordering and translate the expected value of I_t and Q_t into an appropriate order-up-to level S_t. For periods where $E(Q_t)$ is positive an order is placed at the beginning of that period of size $S_t - \sum_{j=1}^{J-1} I_{j,t-1}$. The result of the approaches is a static-dynamic solution, where the order moments are fixed but the order size depends on a fixed S_t and the actual value of $\sum_{j=1}^{J-1} I_{jt}$. For practical reasons one may add a correction for the expected waste till the next review period. Such a policy is not necessarily an optimal strategy but is feasible and practically useful.

An optimal state dependent policy assumes that both the order moment and the order quantity depend on the actual value of the stock levels of all vintages $\sum_{j=1}^{J-1} I_{jt}$. That is an optimal rule is a function $Q_t(I_{j,t-1}, \ldots, I_{J-1,t-1})$. The first studies to such an stock-age dependent rule date back to the early seventies. For an overview see Nahmias [8] and Karaesmen et al. [7]. In these early studies, some analytical results are presented for stylised models that allow for mathematical analysis. Practical results for more realistic models are achieved at the beginning of this century by [2], [4], and [3]. In these papers, an optimal strategy is determined by Stochastic Dynamic Programming (SDP) for both periodic (stationary) and non-stationary problems with a positive lead time of one period under the lost sales assumption. In order to solve large scale problems, aggregation of states is used in that paper. In the current paper, we follow a similar approach that solves large scale problems without aggregation of states. Moreover, we add a service level constraint and restrict ourselves to a finite horizon setting with zero lead time and backlogging of unmet demand.

3 SDP approach

Stochastic Dynamic programming is an appropriate technique to approach (1), as the problem is clearly separable in t. Ingredients of a stochastic dynamic program are the state (and state space), the action (and action space), and a state transition function, next to a contribution function and an objective function. The interested reader is referred to the books [1] and [6] for an introduction in (Stochastic) Dynamic Programming. With respect to the state space, first notice that the waste I_{tJ} does not influence future decisions, as it is not further available; it is not a state variable. The state values are given by $X_t = (I_{1,t}, \ldots, I_{J-1,t})$ in $(J-1)-$dimensional space. In this $(J-1)$-dimensional space, the transition is provided by (4) and (5), so abstractly we have a state transition function Φ:

$$X_t = \Phi(X_{t-1}, Q_t, d_t), \quad t = 1, \ldots, T \tag{6}$$

and we are looking for a rule $Q_t(X_{t-1})$. To facilitate notation, it is convenient to denote the total available inventory at the beginning of period t to fulfil demand in that period by

$$Y_t = \sum_{j=1}^{J-1} X_{j,t-1}.$$

Now the chance constraint including the nonnegativity of Q_t, can be written as

$$Q_t \geq (\Gamma_t^{-1}(\alpha) - Y_t)^+. \tag{7}$$

The waste I_{Jt} is a function of the inventory at the beginning of the period and the demand; $I_{Jt} = f(X_{t-1}, d_t)$. We can write the expected contribution to the objective function in period t as function of state X_{t-1} and decision Q_t:

$$EC(X_{t-1}, Q_t) = g(Q_t) + E\{wf(X_{t-1}, d_t) + h\mathbf{1}^T\Phi(X_{t-1}, Q_t, d_t)\}, \tag{8}$$

where $\mathbf{1}$ is the all-ones vector. Notice that cost and transition functions are not time dependent in the presented model, although the same approach holds for time dependent procurement cost.

The SDP objective function can be written down in a conceptual way via the Bellmann equation using a value function V:

$$V_t(X) = \min_Q \left(EC(X, Q) + E[V_{t+1}(\Phi(X, Q, d_t))]\right), \tag{9}$$

subject to Q fulfilling (7). The argmin of (9) represents the optimum strategy $Q_t(X)$. So starting with a valuation $V_T(X)$ for every possible closing inventory state X, one can compute the valuations backward to know $V_{t-1}(\cdot), \ldots, V_1(\cdot)$ and the optimizing order quantity function $Q_t(\cdot)$. The final result $Q_t(X)$ tells us how much to order in period t given the state of inventory is X. It represents a decision function or a rule that results into the minimum expected cost $V_1(X_0)$ over the time horizon. In a rolling horizon situation, only the optimal decision $Q_1(X_0)$ is implemented and after observing the new inventory levels and updating the demand predictions for the next T periods the SDP may be is executed once again to determine the next order quantity.

Our focus is on how to code the computations of the SDP approach to compute $Q_t(X)$ efficiently. Therefore we first look into the easier characteristics of the case where d_t is deterministic in Section 3.1. In Section 3.2 we go into practical and theoretical considerations for the stochastic SDP approach.

3.1 Solution properties for deterministic demand

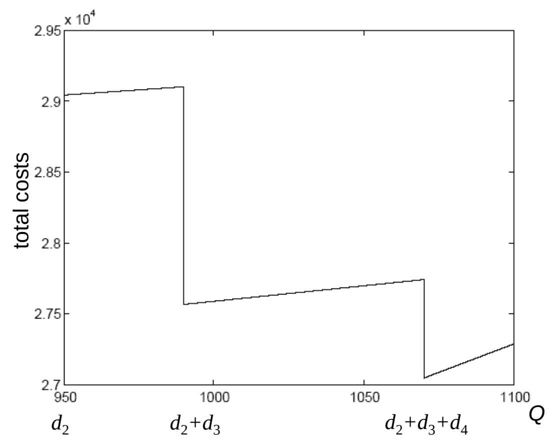

Fig. 1. Minimum total cost at $t = 2$ for $I_{1t} = I_{2t} = 0$ to the end of the planning horizon as function of order decision Q

For the case known as variable demand, where d_t is given for a finite horizon, waste and backlogging can be avoided. As we are dealing with a perishable product, it can be derived that the order Q_t consists of a sum of future demand for an integer number of periods:

$$Q_t \in \{0, d_t, d_t + d_{t+1}, \ldots, d_t + d_{t+1} + \ldots + d_{t+J-1}\}.$$

Notice, that this is optimal in a situation with backlogging. In a lost sales situation, the decision maker can simply order less depending on the interpretation of the service level constraint. Using (deterministic) Dynamic programming (DP), a solution of the Bellmann equation

$$V_t(X) = \min_{Q_t} \left(g(Q_t) + V_{t+1}(\Phi(X, Q_t, d_t)) \right), \quad (10)$$

can easily be found; if demand is larger than the inventory on hand, $Y_t < d_t$, order a sum of future demands and else, do not order at all. Practically, as no analytical expression can be derived, the implementation of the value functions V_t, consists of discretising the state space $\mathcal{X}$ and using interpolation within (10) to evaluate the state one arrives at taking a decision Q_t.

In order to implement a discretisation of the state space, one should have a clue about the range of the state variables X_j. In the deterministic case one can take the range $[0, U_{jt}]$ of X_{jt} and choose upper bound U_{jt} as

$$U_{jt} \geq d_{t+1} + \ldots + d_{t+1+J-j}, \ t = 1, \ldots, T.$$

Implicitly $d_k = 0$ here if $k > T$. We illustrate with the following instance.

Table 1. Varying demand d_t and upper bounds for the inventory, $J = 3$

t	1	2	3	4	5	6	7	8	9	10	11	12
d_t	1900	950	40	80	30	150	800	950	1100	350	150	700
$I_{1t} \leq$	990	120	110	180	950	1750	2050	1450	500	850	700	0
$I_{2t} \leq$	950	40	80	30	150	800	950	1100	350	150	700	0

Example 1. Consider an instance with a perishability of $J = 3$ periods and costs given by $c = 2, h = 1, w = 4$ and $k = 3000$. Table 1 provides the data for the demand d_t and also upper bounds for the inventory if inventory costs and waste are minimised; $I_{1t} \leq d_{t+1} + d_{t+2}$ and $I_{2t} \leq d_{t+1}$. Taking a global upper bound for the inventory implies using $U_1 = \max_t\{d_{t+1} + d_{t+2}\} = 2050$ and $U_2 = \max_{t \geq 2} d_{t+1} = 1100$.

Implementation of the DP approach discretising the state space with steps of 10, leads to the optimal decision sequence that can easily be verified: $Q_1 = 2890, Q_4 = 260, Q_7 = 1750, Q_9 = 1600, Q_{12} = 700$ resulting in a total cost of 32360. In a discretised state space, the DP approach requires solving (10) for each grid point in that space. An illustration is given in Fig. 1, which shows the objective value of (10) for $t = 2$ and $X = (0, 0)$ for varying values of Q starting at d_2. The first slope denotes the inventory cost of one period, the second slope the inventory cost of two periods and the third one represents the increasing cost of waste caused by ordering more than $d_2 + d_3 + d_4$.

The figure shows that in fact one is iteratively minimising a global optimisation problem, see [5]. For the deterministic case, it has easy to determine candidates for the minimum points. As will be illustrated, the stochastic model inherits the global optimisation character, but the minimum points are not that easy to determine.

3.2 Solution properties for stochastic demand

In the stochastic model, the demand is a random variable with known distribution functions Γ_t. The stochastic model allows negative inventory values in 5%

of the cases when a 95% service level applies. One may expect additional costs due to additional production runs when demand appeared to be too big and due to product waste, which may be inevitable.

There are several complications to deal with when we are confronted with stochastic demand. How to deal with the probability distribution of the demand? How to bound the state space? How to solve (9) iteratively?

For notational ease, assume that the demand is Normally distributed with the same coefficient of variation cv over all periods: $d_t \sim \mu_t \times (1 + cv \times N(0,1))$. In this way, demand is fulfilled with a probability of $\alpha\%$,if at the beginning of the period more than $s_t := \Gamma_t^{-1}(\alpha) = (1 + cvG^{-1}(\alpha))\mu_t$ of the product is available, where G is the cdf of the standard normal distribution. For the illustration, Table 2 shows the so-called safety stock s_t and μ_t that corresponds to the data of Example 1 and a $cv = 0.33$, service level $\alpha = 95\%$. The analogy with the deterministic case is that no order is required if the current stock Y_t at the beginning of the period is larger than safety stock s_t.

Table 2. 95% safety level s_t for mean demand μ_t and $cv = 0.33$

t	1	2	3	4	5	6	7	8	9	10	11	12
μ_t	1900	950	40	80	30	150	800	950	1100	350	150	700
s_t	2931	1389	62	123	46	231	1234	1466	1697	540	231	1080

One can discretise the space of possible outcomes of the stochastic demand by using the quantiles of the normal distribution. Practically this works by using an equidistant grid over the probability range $[0, 1]$ with a step p and generating a discrete outcome space $\{\Gamma^{-1}(p), \Gamma^{-1}(2p), \Gamma^{-1}(3p), \ldots, \Gamma^{-1}(1-p)\}$. The consequence of this operation is that the outcome space is truncated by the p-quantiles and every outcome has the same probability of occurrence.

The expected values of the cost and valuation of the state one arrives at, is approximated by an average using the discrete outcomes and the probability. In the deterministic example, the step size in the grid was chosen such that the transition leads to other grid points. An additional complication here is that the transition of all possible outcomes will lead to points in between the grid points requiring the use of interpolation. If the ranges of the state space are not chosen large enough, also extrapolation is required.

The ranges for the stock are less easy to derive than in the deterministic case. Conceptually, there are no bounds, as d_t may not have a bounded support. However, due to the truncation of the outcome space, we have $d_t \in [dmin_t, dmax_t] = [G^{-1}(p) + 1, G^{-1}(1-p) + 1] \times \mu_t \times cv$. As it makes no sense to order more than $Qmax_t = dmax_t + dmax_{t+1} + \ldots + dmax_{t+J-1} - Y_t$, we have an upper bound on the decision Q_t. Using $dmin_t$ as the minimum that will be demanded, we also have bounds on the inventory. For instance $i_{1t} \leq Qmax_t - dmin_t$. The inventory of one period old can never get more negative than $s_t - dmax_t$.

Finally, for every grid point X in the state space, we should now approximate (9) by interpolation of V_t given the discretised values of d_t and find the best order quantity $Q_t(X)$. The difficulty to do so is illustrated in Fig. 2. It shows problem

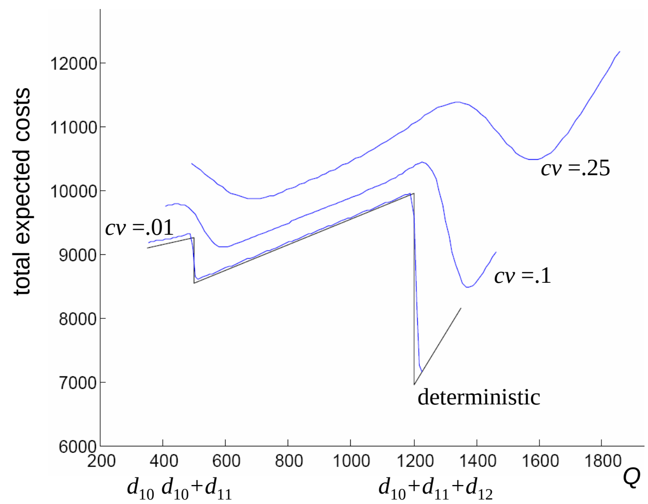

Fig. 2. Minimum total expected cost at $t = 10$, $I_{1t} = I_{2t} = 0$ to the end of the planning horizon as function of order decision Q for different variation coefficient values.

(9) for the last three periods of Example 1. We are again dealing with a one-dimensional global optimization problem for each grid point. Notice that with increasing uncertainty the function to be minimised gets more smooth, leads to higher expected cost and gives the tendency to order for less periods ahead. The figure also illustrates, that when the *cv* goes up from .1 to .25 the best order quantity is no longer in the range of ordering for 2 periods ahead. For the minimization, first the best value for Q was determined over a set of gridpoints within a range of Q. Then from that point, a local search procedure FMINBND of MATLAB was called to finetune the best value.

4 Comparative study

The next question is what is the quality of the described approach in terms of effectiveness and efficiency. With respect to effectiveness, one can estimate the expected cost of the strategy Q_t, by running a simulation with a large number of pseudo randomly generated demand series. In [9], $N = 5000$ series were used to estimate the expected cost with an accuracy of 99%. Moreover, one can test, whether the chosen strategy is really feasible with respect to the chance constraint by keeping track on the number of times that a negative inventory level is reached.

With respect to efficiency, in a decision support environment with a rolling planning horizon, the optimization should be repeated on a daily basis. That means that calculation times should at least be smaller than say an hour in order to fill in new data and place the order. The calculation time depends on the chosen grid density in the state space and is linear in T. It also depends on the goodness of the implementation, the programming language and the platform on which it is run.

We first illustrate the approach with the base case. Then we consider again what happens if variation is going up.

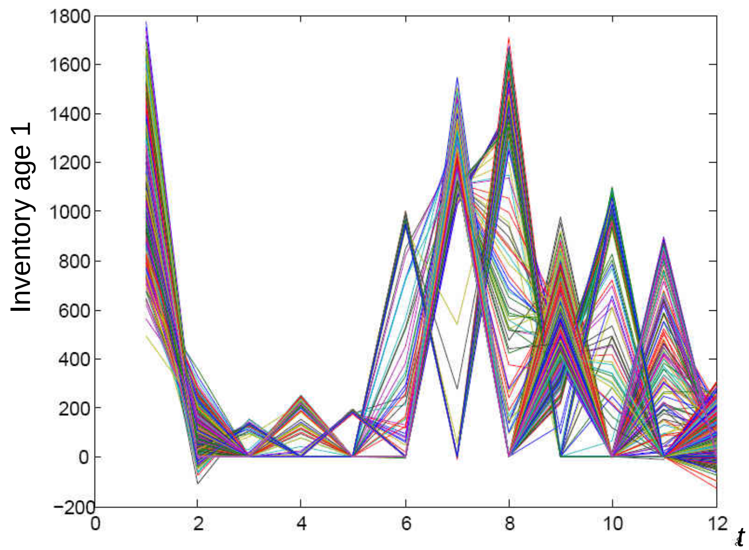

Fig. 3. Development of inventory of age 1 when following the SDP strategy for 2000 generated demand patterns, $cv = 0.1$.

Example 2. Consider again the case of Example 1 with $cv = 0.1$. The SDP procedure was implemented in MATLAB on a standard PC. The state space consist of a grid of $100 \times 60 = 6000$ points for which the value function is determined for $t = 3, \ldots, 12$. For $t = 2$, $X_2 = 0$, so only 100 values are determined and finally $V_1(0) = 3776$ and $Q_1(0) = 3060$ are determined. For the demand d_t, 50 outcomes are generated and evaluated. The calculation of the complete decision table took 500 seconds.

For illustration and evaluation, in total $N = 2000$ demand series were generated for the 12 periods according to the demand distribution $d_t \sim N(\mu_t, 0.1\mu_t))$. The outcome of the SDP procedure Q_t is evaluated for all repetitions providing an average cost of 38126 which is close to the prediction by the SDP of

$V_1(0) = 3776$. As can be observed in Fig. 3, negative inventory levels are only reached at the end of period 2 and at the end of the time horizon. In fact, negative levels are reached at the end of period 2 in 2.2% of the series and in 4.6% of the series at the end of the horizon. That means that chance constraints for the individual periods do not seem binding. Figure 4 provides the development of the inventory of two periods old. One can observe products left at the end of the horizon. However, it is not counted for waste as the items can still be sold in period $T + 1$.

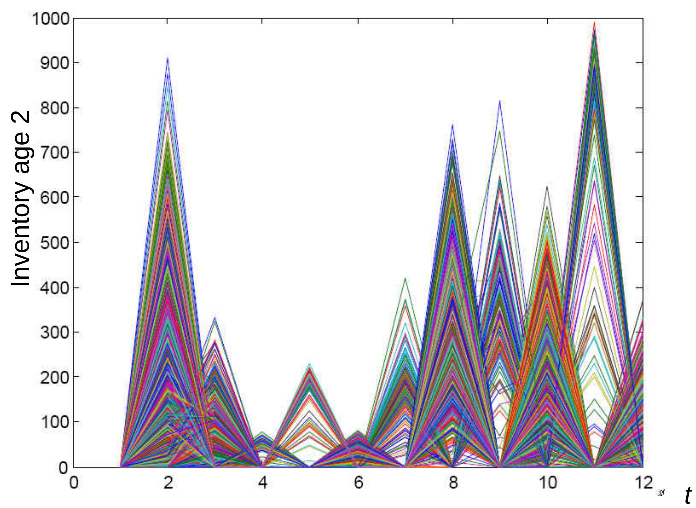

Fig. 4. Development of inventory of age 2 periods when following the SDP strategy for 2000 generated demand patterns, $cv = 0.1$.

We now experiment further with the base case testing the robustness of te algorihm and the behaviour of the model with increasing variation modelled by the coefficient of variation cv. The estimates are based on 5000 repetitions (runs) of a pseudo randomly generated demand series. The results are summarized in Table 3. Over the 5000 repetitions the average order quantities and waste are measured to approximate the resulting mean of the generated order policy $Q_t(X)$. Also the number of runs that lead to out of stock (negative inventory) is also counted for each period in order to check the fulfilment of the service level constraint. In fact, the computational effort of all three generated scenarios is in principle the same, keeping the number of grid points equal.

What can typically be learned from the observations, is that increasing uncertainty leads to ordering for less periods in the optimal strategy. Due to the nonstationarity of the demand, one can see that the critical periods with re-

Table 3. Average values base case over 5000 runs, increasing uncertainty, %oos: percentage of runs that lead to backlog in that period (negative inventory)

	$cv = 0.1$			$cv = 0.25$			$cv = 0.33$		
Cost	38126			43141			46097		
t	Q	waste	%oos	Q	waste	%oos	Q	waste	%oos
1	3060	0	0	2681	0	0	2931	0	0
2	64	0	1.7%	592	0	4.2%	538	0	2.3%
3	8	192	0	15	124	0	13	299	0
4	204	13	0	64	204	0	107	226	0
5	56	1	0	194	3	0	194	2	0
6	66	30	0.7%	106	19	1.2%	98	36	1%
7	1952	0	0.6%	2211	0	0	2341	0	0
8	159	0	0.2%	76	2	1%	74	2	1.4%
9	1213	0	0	976	35	2.3%	924	100	4.4%
10	135	5	0.1%	165	0	0	199	0	0
11	277	71	0.4%	141	121	0	172	149	0
12	457	0	4%	779	0	4.6%	826	1	4.5%

spect to non-negative demand are shifting to other periods when the variation increases. The expected behaviour of more emergency orders (with on average higher fixed order costs), more waste, and higher stock due to higher safety stocks, can also be observed.

5 Conclusions

A description is given of a Stochastic Dynamic programming model to generate finite horizon (T periods) production policies for a perishable (lifetime J) product confronted with a non-stationary demand. The properties of the model and the potential of an SDP approach to come to an optimum strategy are investigated.

In an implementation setting of the concept of dynamic programming, we investigated the boundaries of state space, decision space and the discretisation of the state and outcome space of the random process. The iterative solution of the Bellmann equation for each grid point in the state space, implies solving a one-dimensional Global optimization problem. With a higher variation of demand, this problem becomes more smooth.

The total computing time depends polynomially on the grid density chosen in the $(J-1)-$dimensional state space. However, it increases only linearly in the time horizon T of the problem. For an MILP approximation of the policy as suggested in [9], this is more cumbersome, as the computing time increases in principle exponential in the time horizon due the increase in binary variables. A MATLAB implementation of the code requires in the order of magnitude of minutes and optimal strategy for a given time horizon of $T = 12$ periods, that coincides with the practical problem the research is founded on.

Experimenting with the model and the coefficient of variation that quantifies the uncertainty in demand, one can observe that optimal orders have the tendency to cover less periods if uncertainty goes up. Of course also the cost is going up due to more waste, higher safety stocks and more emergency production orders. The model can be practically used in a rolling horizon setting, where only the first order is carried out and a new plan is generated based on new demand forecasts.

References

1. Bellman, R.: Dynamic Programming. Princeton University Press (1957)
2. Blake, J.T., Thompson, S., Smith, S., Anderson, D., Arellano, R., Bernard, D.: Using dynamic programming to optimize the platelet supply chain in nova scotia. In: M. Dlouhý. Prague, C.R.O. (ed.) Proceedings of the 29th Meeting of the European Working Group on Operational Research Applied to Health Services. pp. 47–65 (2003)
3. Haijema, R., van Dijk, N.M., van der Wal, J., Smit Sibinga, C.: Blood platelet production with breaks: Optimization by SDP and Simulation. International Journal of Production Economics 121, 467–473 (2009), doi:10.1016/j.ijpe.2006.11026
4. Haijema, R., van der Wal, J., van Dijk, N.M.: Blood platelet production: Optimization by dynamic programming and simulation. Computers and Operations Research 34(3), 760–779 (March 2007), doi:10.1016/j.cor.2005.03.023
5. Hendrix, E.M.T., Toth, B.G.: Introduction to Nonlinear and Global Optimization. Springer, New York (2010)
6. Howard, R.A.: Dynamic Programming and Markov Processes. MIT Press, Cambridge University (1960)
7. Karaesmen, I., Scheller-Wolf, A., Deniz, B.: Planning Production and Inventories in the Extended Enterprise, International Series in Operations Research & Management Science, vol. 151, chap. 15. Managing Perishable and Aging Inventories: Review and Future Research Directions, pp. 393–436. Springer (2011)
8. Nahmias, S.: Perishable inventory theory: A review. Operations Research 30, 680–708 (1982)
9. Pauls-Worm, K.G.J., Rossi, R., Haijema, R., Hendrix, E.M.T., van der Vorst, J.G.A.J.: Non-stationary inventory control for a perishable product. OR Spektrum (2012)