

A simple heuristic for perishable item inventory control under non-stationary stochastic demand

Alejandro Gutierrez-Alcoba,^{a*} Roberto Rossi,^b Belen Martin-Barragan,^b Eligius M. T. Hendrix,^a

^a*Computer Architecture, Universidad de Málaga, Málaga, Spain*

^b*Business School, University of Edinburgh, Edinburgh, UK*

(v1.0 released September 2015)

In this paper we study the single-item single-stocking location non-stationary stochastic lot sizing problem for a perishable product. We consider fixed and proportional ordering cost, holding cost, and penalty cost. The item features a limited shelf life, therefore we also take into account a variable cost of disposal. We derive exact analytical expressions to determine the expected value of the inventory of different ages. We also discuss a good approximation for the case in which the shelf-life is limited. To tackle this problem we introduce two new heuristics that extend Silver's heuristic and compare them to an optimal Stochastic Dynamic Programming (SDP) policy in the context of a numerical study. Our results demonstrate the effectiveness of our approach.

Keywords: Stochastic models; Monte Carlo method; Silver's heuristic; perishable inventory control; lot sizing;

1. Introduction

According to the Food and Agriculture Organisation of the United Nations (FAO), around one third of the food that is produced worldwide for human consumption is lost or wasted every year; this amounts to about 1.3 billion tons per year. This loss can be translated into a waste of different valuable resources such as land, water or energy. For this reason, research in perishable item inventory control represents an area of increasing interest.

Lot sizing problems have been long and widely studied due to their important role in economy. Dynamic lot sizing was first introduced by Wagner and Whitin (1958), who discuss a polynomial time exact solution method. Many efficient heuristics appeared over the years in the literature, see e.g. the linear time heuristic introduced in Silver and Meal (1973).

Models considering uncertainty for the demand appeared later in the literature. Silver (1978) presented an heuristic that looks only one cycle ahead and that is based on three different stages: deciding when to order, the number of periods of the cycle and the order quantity. The heuristic introduced in Askin (1981) determines the replenishment levels minimising the incurred cost per period. Bollapragada and Morton (1999) improve the latter both in cost and computational time. Bookbinder and Tan (1988) proposed a heuristic in which first the replenishment periods are decided and after that the quantities are fixed. In Tarim and Kingsman (2006) an MILP (Mixed-Integer Linear Programming) model is used to decide replenishment moments and quantities simultaneously; Rossi, Kilic, and Tarim (2015) generalise Tarim and Kingsman's model to handle a range of service level measures as well as lost sales. Caliskan-Demirag, Chen, and Yang (2012) consider ordering policies for systems in which the fixed cost is dependent on the order size, in a step function for two or multiple values, deriving policies for these cases. Yang, Chen, and Zhou (2014) determine a joint ordering and dynamic pricing strategy for three different models and characterize the optimal policy when inventory cost-rate functions are convex or quasi convex.

This work is an extended version of Rossi (2013)

*Corresponding author.

All aforementioned works operate under a non-stationary demand assumption; the importance of developing models that are able to compute optimal or near-optimal non-stationary policies has been discussed by Tunc et al. (2011).

A review of the early literature on lot sizing for perishable items is provided in Nahmias (1982); it surveys ordering policies from 1960 to 1982. Karaesmen, Scheller-Wolf, and Deniz (2011) make an extensive review of more recent literature for perishable products with fixed or random lifetimes and considering both discrete and continuous models. Bakker, Riezebos, and Teunter (2012) review perishable models since 2001.

In lot sizing for perishable items the structure of the optimal replenishment policy is typically complex: the replenishment quantity depends on the individual age categories of current inventories, as well as on all outstanding orders; for this reason developing effective heuristic policies is of great practical importance in inventory systems for perishable items. According to the above reviews, over the last ten years several inventory models have been derived for controlling perishable item inventory systems. For instance, Minner and Transchel (2010) present a model similar to the one discussed in this paper dealing with a periodic review with service-level constraints; however the role of the fixed ordering cost is not considered in their work. More recently, Hendrix et al. (2012) introduces an SDP approach for a single perishable item subject to non-stationary demand and an α service level constraint. Li and Yu (2014) apply multi-modularity to three dynamic inventory problems; they consider perishability in one of them, for clearance sales, following FIFO (first in, first out) issuance. Chen, Pang, and Pan (2014) analyze a joint pricing and inventory control problem for perishables, considering both backlogging case and lost-sales case; they allow that inventory can be discarded before it is perished. Pauls-Worm et al. (2014) present an MILP approximation model for a YS policy — in this policy Y_t is an indicator variable that is set to one if there is a replenishment up to inventory level S_t in period t — for an inventory control problem under α service level constraints, non-stationary demand and a single item with a fixed shelf life. Alcoba et al. (2015) discuss an YQ policy — in this policy Y_t is an indicator variable that is set to one if there is a replenishment of size Q_t in period t — for a similar setting in which, however, a service level constraint is imposed on lost sales. Pauls-Worm et al. (2015) present an MILP approximation for a YQ policy obtaining costs that are less than 5% higher than those of the optimal policy.

This paper, which extends and complements the discussion of Rossi (2013), focuses on the case of a single perishable item with fixed shelf-life. We consider fixed and variable ordering costs as well as holding and penalty costs for backordered products; the demand is non-stationary over a finite set of periods; inventory is issued following the FIFO policy. A major difference with respect to Rossi (2013) is the adoption of a penalty cost scheme in place of a service level constraint; furthermore, we present a broader set of instances and we assessed the quality of our new heuristics against optimal solutions obtained via stochastic dynamic programming.

We make the following contribution to the literature on perishable item non-stationary stochastic lot sizing.

- We introduce exact analytical expressions to compute the expected value of the inventory for different product ages when a product can age indefinitely; the expressions work under discrete as well as continuous demand distributions.
- We derive an analytical approximation for the case in which the product age is discrete and finite.
- By using the former results, we extend Silver's heuristic (Silver 1978) to the case in which the product is perishable; in particular we introduce analytical and simulation-based variants of the approach.
- We conduct an extensive numerical study which demonstrates that our new heuristics lead to a cost that, on average, is 5% higher than the optimal cost.

The rest of this work is structured as follows. Section 2 describes the problem setting and the stochastic dynamic programming model. Section 3 discusses various theoretical results that are then employed, in Section 4, to develop an analytical extension of Silver's heuristic; moreover, in the same section we develop a more accurate but computationally less efficient heuristic approach. A computational study and its results are detailed in Section 5; finally, Section 6 draws conclusions.

2. Problem and model description

We consider a single-item, single-stock location, production planning problem over a time horizon of T periods. The product we consider is perishable and its age, in periods, is denoted by $a \in \{1, \dots, A\}$; $a = 1$ denotes fresh products delivered at the beginning of the current period. At the end of a given period t , all items of age A are scrapped; we will consider $A < T$ to ensure that perishability has an effect in the model.

The stochastic demand we consider is non-stationary, which means its distribution varies from one period to the next. Demand in period t is represented by means of a non-negative random variable D_t with known cumulative distribution function F_t . Random demand is assumed to be independent over the periods. The system operates under a FIFO product issuing policy. Unmet demand at any given period is backordered and satisfied as soon as a replenishment arrives. The sellback of excess stock is not allowed.

For ease of exposition, we assume that delivery of products is instantaneous, i.e. lead time is zero; both our heuristics can be extended to situations in which lead-time is positive and deterministic, this extension will be briefly discussed in the Appendix. There is a fixed ordering cost o for placing a replenishment and a variable ordering cost v proportional to the order quantity; a variable inventory holding cost h is charged on every product unit carried from one period to the next, regardless of its age; a variable penalty cost p is charged for each back-ordered unit at the end of each period during which it remains unsatisfied; a variable wastage cost w is charged for every item of age A discarded at the end of a period. Our aim is to find a replenishment plan that minimizes the expected total cost, which is composed of ordering costs, holding costs, penalty and waste costs over a T -period planning horizon.

Let the net inventory denote the on-hand stock minus backorders. The model assumes that the sequence of events in each period is as follows. At the beginning of a period the decision maker checks the net inventory for different product ages that has been carried over from the former period. Since we are dealing with complete backlogging, this quantity may be negative. However, note that only fresh products can be backordered, since the supplier only delivers fresh products. On the basis of this information, if necessary, the decision maker issues an order. Then the random demand is observed and items are issued to meet demand according to a FIFO policy. After demand is realised at the end of the period inventory ages and if the net inventory is positive, any item of age A that remains in stock is discarded and waste cost is incurred. Items of age 1 to $A - 1$ that remain in stock are carried over to the next period and holding costs are incurred. If the net inventory is negative, penalty cost is incurred on any unit backordered.

For convenience, the notation adopted in the rest of this work is summarized in Table 2. The problem can be naturally modeled as a stochastic dynamic problem by considering the following components.

States. The system state at the beginning of period t is represented by the vector $\mathbf{I}_{t-1} = (I_{t-1}^1, \dots, I_{t-1}^{A-1})$; note that this is also the system state at the end of period $t - 1$, and that the waste W_{t-1} (i.e. I_{t-1}^A) does not determine future decisions; inventory of age 1 is the only value in the state vector that can take negative values to denote state of the cumulated backorders.

Actions. The action is the order quantity Q_t for period t given initial inventory \mathbf{I}_{t-1} .

State Transition Function. The state transition follows FIFO dynamics. The inventory at the end of period t is determined by the inventory from the previous period, the realised demand D_t and the delivery of the ordered quantity Q_t . The notation $(\cdot)^+ = \max(\cdot, 0)$ is used to shorten the expressions. Given state \mathbf{I}_{t-1} and action Q_t , the state transition of the system is described by the following expressions. Waste, i.e. items of age A , at the end of period t is defined as

$$W_t = (I_{t-1}^{A-1} - D_t)^+, \quad t = 1, \dots, T. \quad (1)$$

Indices	
t	period index, $t = 1, \dots, T$, with T the time horizon
a	age index, $a = 1, \dots, A$, with A the fixed shelf life
Problem parameters	
o	fixed ordering cost
v	unit procurement cost
h	unit inventory cost
p	unit penalty cost
w	unit disposal cost
i^a	initial inventory of age a at the beginning of the planning horizon
Random variables	
D_t	random demand in period t with cumulative distribution function F_t
State variables	
\mathbf{I}_t	system state (i.e. inventory vector) at end of period t
I_t^a	inventory of age a at end of period t , where $I_t^a \geq 0$ for $a = 2, \dots, A$
W_t	waste at the end of period t (i.e. I_t^A)
Decision variables	
Q_t	order quantity at the beginning of period t given inventory $I_{t-1}^1, \dots, I_{t-1}^{A-1}$

Table 1. Notation adopted in this work.

Inventory of age $a \in \{2, \dots, A-1\}$ that can still be used in the next period follows from

$$I_t^a = \left(I_{t-1}^{a-1} - (D_t - \sum_{j=a}^{A-1} I_{t-1}^j)^+ \right)^+, \quad t = 1, \dots, T; \quad a = 2, \dots, A-1. \quad (2)$$

Finally, inventory of fresh products/backorders follows from

$$I_t^1 = Q_t - (D_t - \sum_{a=1}^{A-1} I_{t-1}^a)^+, \quad t = 1, \dots, T. \quad (3)$$

Immediate Costs. Given state \mathbf{I}_{t-1} and the described transformation towards state \mathbf{I}_t as a consequence of action Q_t , the immediate costs are given by

$$C(\mathbf{I}_{t-1}, Q_t) = g(Q_t) + E \left(h \sum_{a=1}^{A-1} (I_t^a)^+ + p(-I_t^1)^+ + w W_t \right), \quad (4)$$

where E denotes the expectation taken with respect to random demand D_t and function g is defined as:

$$g(Q) = \begin{cases} o + vQ & \text{if } Q > 0 \\ 0 & \text{if } Q = 0. \end{cases} \quad (5)$$

Objective Function. The objective is to find an order policy that minimizes the expected total cost over the T period planning horizon, that is

$$\min_{Q_1} \left(C(\mathbf{I}_0, Q_1) + \min_{Q_2} \left(C(\mathbf{I}_1, Q_2) + \dots + \min_{Q_T} C(\mathbf{I}_{T-1}, Q_T) \right) \right). \quad (6)$$

3. Perishable stock dynamics: theoretical results

The purpose of this Section is to provide exact analytical expressions for the expected inventory level at different ages for the case where items can age indefinitely. These results are used to provide a good approximation of the expected value of the inventory levels for the case where items perish after A periods.

Consider the single-item single-stocking location inventory system over a single period t subject to a random demand D with known cumulative distribution F , probability density function f , and expected value μ . Since we refer to a single period problem, in order to keep notation compact, we have dropped the index t from D_t , F_t , f_t and μ_t .

We have seen that only I_t^1 can take a negative value, whereas older age inventory is bound to be nonnegative. For convenience, we introduce $I_{t-1}^0 = Q_t$, so that we can denote the overall inventory available at the beginning of the period as

$$Y = \sum_{a=0}^{A-1} I_{t-1}^a;$$

we also exploit notation I_t^A to denote the waste W_t .

Our aim is to find analytical expressions for $E(I_t^a)$, i.e. the expected value of I_t^a , for $a = 1, \dots, A$. The following Lemmas provide analytical expressions to calculate the expected value of inventory levels of different ages, when they can age indefinitely, both for discrete and continuous demand support.

Lemma 1. *Let D be a random variable defined over a continuous support and $Y = \sum_{a=0}^{A-1} I_{t-1}^a$,*

$$E(-I_t^1)^+ = \int_0^Y (Y - x)f(x) dx - (Y - \mu).$$

Proof. See Rossi et al. (2014), Lemma 3 □

Lemma 2. *Let D be a random variable defined over a discrete support $\mathbb{S} \subseteq \mathbb{N}$ with a positive probability mass function $f(x) = F(x) - F(x-1)$ for $x \in \mathbb{S}$ and zero elsewhere. Let $Y_a = \sum_{j=a-1}^{A-1} I_{t-1}^j$, $Y_{a+1} = \sum_{j=a}^{A-1} I_{t-1}^j$, and $Y_{A+1} = 0$, then for $a = 1, \dots, A$*

$$E(I_t^a)^+ = \sum_{x=0}^{Y_a-1} F(x) - \sum_{x=0}^{Y_{a+1}-1} F(x).$$

Proof. Note that $I_{t-1}^{a-1} = Y_a - Y_{a+1}$; moreover, for any value $Y \in \mathbb{S}$ we have $\sum_{x=0}^Y (Y - x)f(x) = \sum_{x=0}^{Y-1} F(x)$ (Rossi et al. 2014). Following the introduced symbols and inventory dynamics of

Equations (1), (2) and (3), we have

$$\begin{aligned}
 E(I_t^a)^+ &= \sum_{x=0}^{Y_a} (I_{t-1}^{a-1} - (x - Y_{a+1})^+) f(x) \\
 &= \sum_{x=0}^{Y_a} (Y_a - Y_{a+1} - (x - Y_{a+1})^+) f(x) \\
 &= \sum_{x=0}^{Y_{a+1}} (Y_a - Y_{a+1}) f(x) + \sum_{x=Y_{a+1}+1}^{Y_a} (Y_a - x) f(x) \\
 &= \sum_{x=0}^{Y_{a+1}} (Y_a - Y_{a+1}) f(x) + \sum_{x=0}^{Y_a} (Y_a - x) f(x) - \sum_{x=0}^{Y_{a+1}} (Y_a - x) f(x) \\
 &= \sum_{x=0}^{Y_a} (Y_a - x) f(x) - \sum_{x=0}^{Y_{a+1}} (Y_{a+1} - x) f(x) \\
 &= \sum_{x=0}^{Y_a-1} F(x) - \sum_{x=0}^{Y_{a+1}-1} F(x).
 \end{aligned}$$

□

Lemma 3. Let D be a random variable defined over a continuous support, $Y_a = \sum_{j=a-1}^{A-1} I_{t-1}^j$, $Y_{a+1} = \sum_{j=a}^{A-1} I_{t-1}^j$, and $Y_{A+1} = 0$, then for $i = 1, \dots, A$

$$E(I_t^a)^+ = \int_0^{Y_a} (Y_a - x) f(x) dx - \int_0^{Y_{a+1}} (Y_{a+1} - x) f(x) dx.$$

Proof. Follows from $I_{t-1}^{a-1} = Y_a - Y_{a+1}$ and a similar reasoning to the one presented in Lemma 2. □

Lemma 4. $E(I_t^a) = E(I_t^a)^+ - E(-I_t^a)^+$.

Proof. Follows from the definition of expectation. □

Lemma 5. For $a = 2, \dots, A$,

$$E(I_t^a) = E(I_t^a)^+.$$

Proof. Follows from the fact that only fresh products can be back-ordered. □

We now consider a horizon comprising T time periods with random demand D_t in each period $t = 1, \dots, T$. Items carried over from one period to the next age and become one time period older. We first analyse the case where *items can age indefinitely*. Our aim is to compute $E(I_t^a)$, i.e. the expected value of I_t^a , for $t = 1, \dots, T$ and $a = t, \dots, A + a - 1$. This question is easy to handle if we reduce a t -period problem to a single-period problem subject to cumulative demand $D = D_1 + \dots + D_t$. Consider $E(I_t^a)$ for this new single-period problem, as previously obtained, and note that the value obtained corresponds to $E(I_t^{t+a-1})$ for the t -period problem.

Example. Consider a $T = 2$ periods planning horizon. Demand D_t in each period t follows a Poisson distribution with rate $\lambda_t = 50$. The initial inventory is $\mathbf{I}_{t-1} = (I_{t-1}^0, I_{t-1}^1, I_{t-1}^2) = (25, 50, 50)$. This means there are 25 fresh items that have just been produced; 50 items that were produced a period before and 50 items that were produced two periods before, which are still in stock at the beginning of the planning horizon. To determine $E(I_1^a)$, $a = 1, \dots, A$ we apply the results just presented, in particular Lemma 2, which lead to $E(I_1) = (25, 47.18, 2.81)$. The value

of $E(I_2^a)$ for $a = 2, \dots, A + 1$ can be derived considering the 2-period problem as a single-period problem subject to a Poisson demand with rate $\lambda = 100$ and we apply once more the results just presented and Lemma 2. This leads to the vector $E(\mathbf{I}_2) = (0, 21.04, 3.98)$. Since we do not place an order in period 2, the expected number of fresh items at the end of period 2 is zero. Items of age 3 at the end of period 1 have aged.

The case in which items of age A — i.e. items that were produced $A - 1$ periods before — are discarded at the end of a period, complicates the exact analytical derivation of $E(I_t^a)$ for $t = 1, \dots, T$ and $a = t, \dots, A + t - 1$. We shall therefore proceed in a pragmatic manner, by introducing a simple and yet powerful approximation that reuses as much as possible results developed so far. We will demonstrate the effectiveness of this approximation in our computational study.

The key intuition leading to our approximation is an inductive argument. Our base case is the determination of $E(I_t^a)$, which can be carried out analytically for $t = 1$ and $a = t, \dots, A$ (with $I_t^A = W_t$) by using the results presented so far. In principle it is also possible to determine analytically the higher moments of the distribution of I_t^a , e.g. the variance $\text{Var}(I_t^a)$ for $t = 1$. To approximate $E(I_t^a)$, we first operate as if items can age indefinitely, i.e. they are never discarded, and we reduce the t -period problem to a single period problem with demand $D = D_1 + \dots + D_t + W_1 + \dots + W_{t-1}$, where distributions of W_1, \dots, W_{t-1} are derived at previous induction steps. Once this new demand distribution is obtained, values for $E(I_j^a)$ for $j < t$ and $a = t, \dots, A - 1$ can be computed iteratively by reusing results developed so far. Once more, one should note that the value $E(I_t^a)$ obtained for this new single-period problem corresponds to $E(I_t^{t+a-1})$ for the t -period problem; of course, $E(I_t^a) = 0$ for $a > A$.

Example. Consider a problem in which D_t follows a Poisson distribution with rate λ_t for $t = 1, \dots, T$ and the shelf life is $A = 3$. The approach is based on an approximation of the distribution of I_t^a by fitting an appropriate first moment (i.e. mean) to a Poisson distribution.¹ This means that in period t , after having carried out the induction over periods $1, \dots, t - 1$, we will reduce the t -period problem to a single period problem subject to Poisson demand with expected value $E(D_1) + \dots + E(D_t) + E(W_1) + \dots + E(W_{t-1})$. More specifically, we refer to the same problem analysed in the previous example. However, now the shelf life is limited to $A = 3$. We first compute $E(\mathbf{I}_1) = (25, 47.18, 2.81)$, for which the computation is identical to the one carried out under the previous example setting. By exploiting this information, and in particular $E(I_1^3) = 2.81$, we reduce the two-period problem to a single period problem under Poisson demand with rate $\lambda = 50 + 50 + 2.81$ and exploit once more Lemma 2 to compute the approximation $(0, 19.47, 2.77)$. This vector is close to the actual vector $E(\mathbf{I}_2) = (0, 20.219, 1.993)$ which can be found carrying out an exact convolution.

If the demand distribution is not uniquely determined by its mean, higher order moments can be obtained and matched in a similar fashion, for instance to a normal distribution.

4. Two new heuristics

In this section, we present two heuristics that exploit the mathematical properties just presented: an analytical extension of Silver's heuristic and a simulation-optimisation heuristic.

An analytical extension of Silver's heuristic. Silver's heuristic (Silver 1978) extends the well-known Silver-Meal lot sizing heuristic (Silver and Meal 1973) to a probabilistic setting. The key idea behind the Silver-Meal lot sizing heuristic is to determine the average cost per period as a function of the number of periods that can be covered by the current order. Exploiting the fact that this function is convex with respect to the number of periods, it is possible to determine the

¹In general, one may want to use more advanced distribution fitting techniques.

Algorithm 1: Analytical extension of Silver's heuristic

Data: the current period t ; initial inventory \mathbf{I}_{t-1} at the beginning of period t
Result: the optimal order quantity, $Q^* \geq 0$

```

1  $r \leftarrow t$ ;
2  $c^* \leftarrow \infty$ ;
3 do
4    $Q \leftarrow \operatorname{argmin}_{Q \geq 0} [C(\mathbf{I}_{t-1}, Q) + C(\mathbf{I}_t, 0) + \dots + C(\mathbf{I}_{r-1}, 0)]$ ;
5    $c \leftarrow (C(\mathbf{I}_{t-1}, Q) + C(\mathbf{I}_t, 0) + \dots + C(\mathbf{I}_{r-1}, 0)) / (r - t + 1)$ ;
6   if  $c \leq c^*$  then
7      $c^* \leftarrow c$ ;
8      $Q^* \leftarrow Q$ ;
9   end
10   $r \leftarrow r + 1$ ;
11 while  $c \leq c^*$  and  $r < t + A$ ;
```

optimal length of the next replenishment cycle. Next, we briefly illustrate how the results presented in Section 3 can be used to extend Silver's heuristic to the case in which inventory comprises items of different age classes. We shall first formally define the concept of *replenishment cycle*.

Definition 1. A *replenishment cycle* (t, r) denotes the time interval between two consecutive replenishments executed at period t and at period $r + 1$.

Consider a replenishment cycle (t, r) where order quantity Q_t aims to cover demand of periods t, \dots, r . The following lemma extends findings in Federgruen and Wang (2015).

Lemma 6. The expected total cost defined by equations (1)...(4) during replenishment cycle (t, r) given starting inventory \mathbf{I}_{t-1} is convex in Q_t for $Q_t > 0$.

Proof. No order is placed in periods $t + 1, \dots, r$, the total costs of the replenishment cycle is

$$C(\mathbf{I}_{t-1}, Q_t) + C(\mathbf{I}_t, 0) + \dots + C(\mathbf{I}_{r-1}, 0).$$

This function is separable, and its separable components are all increasing or decreasing convex piecewise linear functions of Q_t : the unit ordering cost, holding and waste costs increase with Q_t , the penalty cost decreases with Q_t ; this function is therefore convex in Q_t . \square

Our heuristic exploits the results in Section 3 to compute, for a given starting inventory \mathbf{I}_{t-1} , the inventory levels $E(\mathbf{I}_t)$ during the replenishment cycle $t \in \{t, \dots, r\}$. Using Lemma 6 it computes the optimal order quantity Q_t for replenishment cycle (t, r) as well as the associated expected total cost per period. As in Silver's heuristic, we increase the value of r , starting from t , until the expected total cost per period associated with replenishment cycle (t, r) first increases. Let $r + 1$ be such value. The optimal action in period t is to order a quantity Q_t which minimizes the expected total cost for replenishment cycle (t, r) . A pseudocode of the algorithm is given in Algorithm 1. We implement this heuristic in a rolling horizon framework: if the planning horizon comprises T periods, at the beginning of each period $t = 1, \dots, T$ we run the procedure described in Algorithm 1 to determine if an order must be issued and, if so, the respective order quantity. It should be noted that Algorithm 1 is equivalent to Silver's heuristic provided $A = \infty$ and $w = 0$.

Example. Consider an instance over a planning horizon of $T = 3$ periods, with shelf life $A = 3$. Demand in each period t is Poisson distributed with rate λ_t , where $\lambda \in \{4, 3, 3\}$. Fixed ordering cost o is set to 10, holding cost is set to 1, waste cost w is set to 2, penalty cost p is set to 5. Initial inventory is $\mathbf{I}_0 = (I_0^1, I_0^2) = (1, 1)$. We consider replenishment cycles (t, r) of increasing length:

- for $(t, r) = (1, 1)$ the optimal action is to not order ($Q = 0$) and use existing inventory to cover demand in period 1, the expected total cost per period of this policy is 10.67 — note

Algorithm 2: Simulation-optimisation heuristic

Data: the current period t ; initial inventory \mathbf{I}_{t-1} at the beginning of period t
Result: the optimal order quantity, $Q^* \geq 0$

- 1 $r \leftarrow t$;
- 2 $c^* \leftarrow \infty$;
- 3 **do**
- 4 $Q \leftarrow \operatorname{argmin}_{Q \geq 0} [\widehat{C}_{t,r}^N(\mathbf{I}_{t-1}, Q)]$;
- 5 $c \leftarrow (\widehat{C}_{t,r}^N(\mathbf{I}_{t-1}, Q)) / (r - t - 1)$;
- 6 **if** $c \leq c^*$ **then**
- 7 $c^* \leftarrow c$;
- 8 $Q^* \leftarrow Q$;
- 9 **end**
- 10 $r \leftarrow r + 1$;
- 11 **while** $c \leq c^*$ and $r < t + A$;

that, if we decided to place an order, the optimal order quantity to cover demand in period 1 would be $Q = 3.96$ and the expected total cost per period would increase to 13.21;

- for $(t, r) = (1, 2)$ the optimal action is to place an order $Q = 6.04$, the expected total cost per period of this policy is 9.56, since this is less than 10.67 we proceed and consider the next possible replenishment cycle length;
- for $(t, r) = (1, 3)$ the optimal action is to order $Q = 7.99$, the expected total cost per period of this policy is 9.68; since this is greater than 9.56, we conclude that the optimal action in period 1 is to order $Q = 6.04$ and incur an expected total cost per period of 9.56. After observing actual demand in period 1, the procedure can be iterated to determine order quantities in following periods.

Simulation-optimisation heuristic. The approach discussed so far extends Silver's heuristic by using an analytical approximation for the expected value of the inventory levels (Section 3). The second approach uses Monte Carlo simulation to approximate $E(\mathbf{I}_t)$ in a replenishment cycle. Using this sample-based approximation we can estimate the expected total cost of replenishment cycles and therefore develop a simulation-based extension of Silver's heuristic. The heuristic obtained is independent of the specific demand distribution under scrutiny. More precisely, given a cycle (t, r) , the method draws N independent samples of the demand for each period $k = t, \dots, r$. It then generates sample average estimates $\widehat{E}(\mathbf{I}_k)$ of $E(\mathbf{I}_k)$, for $k = t, \dots, r$, as a function of the order quantity Q and of the initial inventory \mathbf{I}_{t-1} , by relying on the inventory dynamics described in Eq. (1), (2) and (3). These estimates $\widehat{E}(\mathbf{I}_k)$ can be immediately translated into a Monte Carlo estimate $\widehat{C}_{t,r}^N(\mathbf{I}_{t-1}, Q)$ of the expected total cost for cycle (t, r)

$$\widehat{C}_{t,r}^N(\mathbf{I}_{t-1}, Q) = \sum_{k=t}^r \left(h \widehat{E}(I_k^1)^+ + h \sum_{a=2}^{A-1} \widehat{E}(I_k^a) + p \widehat{E}(-I_k^1)^+ + w \widehat{E}(I_k^A) \right).$$

This cost estimate can be readily used in the context of an algorithm (Algorithm 2) that mimics the one previously presented for the computation of the optimal order quantity Q^* .

Finally, it is worth sparing some words on the computational complexity of the two approaches. Both our heuristics are extension of Silver's heuristic and computationally speaking they are similar to it. This can be clearly observed by contrasting Algorithms 1 and 2. However, while Algorithm 1, by leveraging on our analytical approximations, features complexity $O(AT \log Q)$; i.e. in the worst case, for each period $t = 1, \dots, T$, we need to minimize a convex function involving at most A age categories over domain $0, \dots, Q$, where $Q = \epsilon \bar{Q}$, ϵ denotes the selected error threshold (i.e. discretization step), and \bar{Q} denotes an upper bound for the order quantity; Algorithm 2 relies on

a sample-based approach exploiting N samples to compute the optimal order quantity at each period $t = 1, \dots, T$, therefore its complexity only increases linearly in N and is $O(ATN \log \mathcal{Q})$.

5. Computational study

In this section we present a computational study that investigates the effectiveness of our heuristics. Section 5.1 outlines the experimental design; Section 5.2 presents the cost differences reached by the heuristics; and Section 5.3 discusses our findings.

5.1 *Experimental design*

To study the goodness of the heuristics presented in Section 4, we follow the experimental setup originally presented in Rossi, Kilic, and Tarim (2015) in which the authors employ a test bed comprising ten different stochastic demand patterns over a time horizon of $T = 15$ periods and a shelf life of the items of $A = 3$ periods. The demand patterns include a stationary demand instance (STA), two life cycle instances (LCY1 and LCY2), two sinusoidal (SIN1 and SIN2), and four empirical (EMP1, ..., EMP4) patterns of expected demand; these patterns are displayed in Figure 1 — full numerical details are provided in Appendix II. The stochastic demand in each period t follows a Poisson distribution mean λ_t ; to ensure that the state space is finite in the SDP implementation we considered the cumulative distribution of the demand over the whole horizon and we limited the support to the values between 0 and the upper 99% quantile of this distribution.

The values of the parameters of the objective function (Eq. 6) have been varied in a systematic way. Note that the unit procurement cost is directly related to the penalty cost p and disposal cost w . Variation of all of them may lead to a colinear design. Moreover, the inventory holding cost h and setup cost o define the length of the replenishment cycle, so only one of them has to be varied. Therefore the value of the procurement cost has been fixed to $v = 0$ while the unit inventory cost is set to $h = 1$. For penalty and disposal costs, values $p \in \{2, 5, 10\}$ and $w \in \{2, 5, 10\}$ have been considered. Finally, the fixed ordering cost is set to the sum of mean demand of all the periods of the instance multiplied by a scale factor $l_o \in \{1, 2.5, 5\}$. The Monte Carlo approximation of the order quantity (Algorithm 2) is based on $N = 300$ samples. To estimate the cost of our heuristics we rely on 500 rolling horizon simulation runs. The full Cartesian space of the different ten demand patterns and the different values for penalty, wastage and fixed costs results into a set of 270 scenarios. The experimental design is based on randomly selecting 20% of them to test the extension of Silver’s heuristic against the optimal solution provided by the stochastic dynamic programming model. Table 4 shows the parameter values of the scenarios of the experimental design. Appendix II reports the 54 instances systematically selected in the context of our study.

5.2 *Results*

Table 2 summarises the results obtained for our computational study when pivoting on different parameters of the test bed: the problem instance, the order cost level $l_o \in \{1, 2.5, 5\}$, the penalty cost $p \in \{2, 5, 10\}$ and the waste cost $w \in \{2, 5, 10\}$. The first column of Table 2 (heading “Analytical”) shows the MPE (Mean Percentage Error) from the optimal cost obtained by an exact SDP model to the cost obtained using our analytical extension of Silver’s heuristic and its 95% CI (Confidence Interval), while the second column (heading “Simulation-optimization”) shows the same figures using our Monte Carlo simulation-optimisation heuristic.

5.3 *Discussion of results*

From Table 2 follows that the simulation-optimization heuristic performs, in general, better than the analytical approximation. However, the latter runs about 10 times faster than the first. On average, the analytical heuristic reaches a cost that is 5.96% higher than the optimal cost for the 54 tested instances, while for the simulation-optimization heuristic the cost difference reduces to

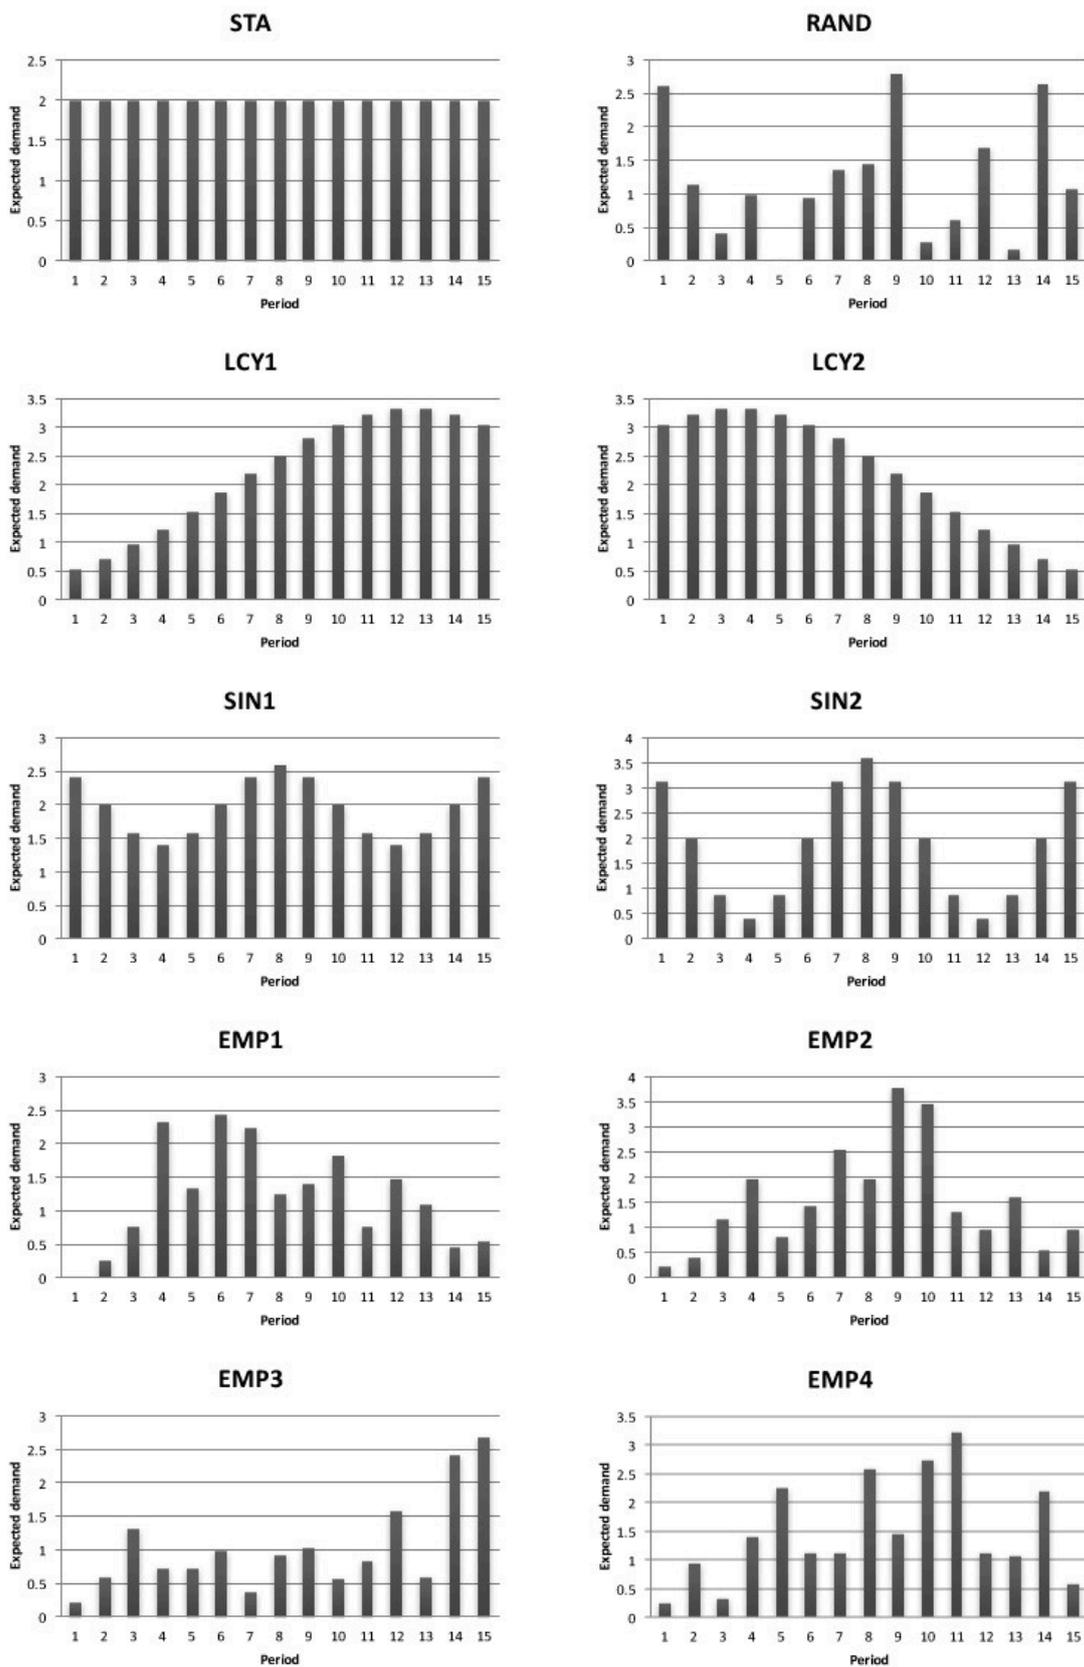


Figure 1. Demand patterns used in our computational study; the values denote $E[D_t]$ in each period $t \in \{1, \dots, T\}$

Table 2. Pivot table for the computational study

	Analytical		Simulation-optimization		Observations
	MPE	0.95 CI	MPE	0.95 CI	
Ord cost level					
1	3.06	±1.77	2.91	± 1.66	18
2,5	5.24	±3.70	4.22	± 2.33	18
5	9.59	±6.36	7.14	± 3.84	18
Penalty cost					
2	2.37	±1.17	2.00	± 0.92	18
5	4.63	±3.04	4.63	± 2.78	18
10	10.8	±6.51	7.65	± 3.56	18
Waste					
1	1.74	± 0.59	2,13	± 0.75	18
5	3.68	± 2.03	3.44	± 1.60	18
10	12.4	± 6.45	8.71	± 3.99	18
Demand pattern					
EMP1	8.62	± 7.33	6.47	± 5.13	10
EMP2	3.53	±4.97	3.13	± 4.83	5
EMP3	8.78	±10.1	7.34	± 5.78	9
EMP4	23.8	±205	11.9	± 102	2
LCY1	2.42	±3.44	2.55	±3.27	5
LCY2	1.06	± 2.32	1.32	± 3.80	3
RAND	14.6	±12.1	12.9	± 21.1	2
SIN1	2.58	± 2.74	3.32	± 3.71	5
SIN2	3.20	± 2.97	2.91	± 2.31	5
STA	2.25	± 2.84	1.95	± 2.75	8
General	5.96	± 2.47	4.76	± 1.57	54
Time (aprox.)	5 secs	50 secs			

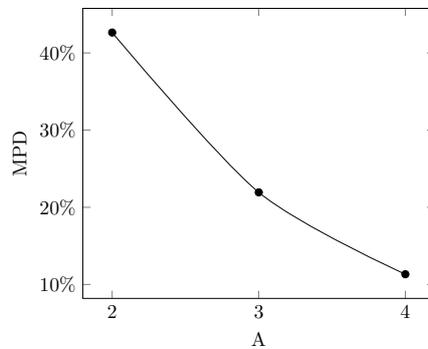


Figure 2. Mean Percentage Difference (MPD) of the optimal solutions of the test bed (Table 4) for A = 2, 3, 4 compared to the non-perishable case

4.76%. It can be observed that when any of the pivoting parameters is larger in comparison to the remaining ones, the heuristics perform worse. For example, for the order cost level, the simulation-optimization heuristic gives 2.91% of difference with the optimal cost for a level equal to 1, while that percentage is 7.144% for level 5. When considering the demand patterns, the cost difference is generally lower than 4%. However, the random pattern (RAND) and empirical patterns (EMP1, EMP3, and EMP4) produce larger differences. In particular, the largest difference, which amounts to 23.8% for the analytical heuristic and 11.9% for the simulation-optimization heuristic is observed for EMP4. The reader should nevertheless note that there are only two observations of this pattern and the associated confidence interval is therefore very large.

By using our exact stochastic dynamic programming approach, we also computed the optimal policy for the instances of the test bed discussed in Section 5.1 when A = 2, 3, 4, and compared their costs to the non-perishable case. As it can be seen in Figure 2, the difference to the non-

perishable case quickly decreases as A increases. This means that for $A > 4$, the advantage of adopting a model that accounts for perishability quickly becomes negligible.

6. Conclusions

In this paper, we presented an inventory control model for perishable items with a fixed shelf-life. After discussing a number of analytical results concerning the dynamics of perishable stock in the inventory system under scrutiny, we introduced two extensions of Silver's heuristic: an analytical approximation and a Monte Carlo simulation-optimisation approach. We demonstrated the effectiveness of these heuristics in a comprehensive numerical study that focused on a number of different demand patterns, as well as on a range of different values for the ordering cost, the penalty cost and the waste cost. Our results show that the simulation-optimisation approach features a better cost performance that, on average, is 5% above the optimal cost, the respective figure for our analytical approximation is 6% .

Acknowledgements

This paper has been supported by The Spanish Ministry (TIN2015-66680-c2-2-R).

Appendix

Appendix I discusses an extension of our heuristics to the case in which lead time is deterministic and positive and Appendix II provides further details on our computational study.

Appendix I: extension to the case in which lead time is deterministic and positive

We refer to a positive and deterministic leadtime L . For both the heuristics discussed in Section 4 the extension is relatively simple and proceeds as follows.

First, it is necessary to clarify if items are fresh when they are shipped or when they are delivered, but this assumption does not substantially affect our reasoning.

The inventory vector \mathbf{I}_t must be expanded to accommodate L additional stock categories, so the state dimension increases, but it increases linearly in L . If shelf life is A the vector now features $A + L$ possible stock categories, where categories $1, \dots, L$ now denote items that will be delivered in $L, L-1, \dots, 1$ periods, respectively. Categories $L+1, \dots, L+A$ denote fresh items just delivered ($L+1$), items that are one period old ($L+2$), and so on and so forth.

When both heuristics try to determine the order quantity at period t , the initial inventory level in the expressions will have to be replaced by the inventory position, i.e. on hand stock plus pending orders not yet delivered, minus backorders.

Computation of the expected cost for the replenishment cycle starting at period t and ending at period $t+k$ will have to take into account the fact that the next order placed in period $t+k+1$ will only be delivered in period $t+k+L$; in practice, this amounts to extending the replenishment cycle by L periods and thus taking into account demand over lead time and associated expected costs (holding, penalty, etc) in periods $t+k+1, \dots, t+k+L$. Depending on how we define our holding cost accounting — i.e. if we should charge or not holding cost on pending orders not yet delivered — holding cost may or may not be charged on stock categories $1, \dots, L$. In any case, since the cost function is separable (Lemma 6) and we are able to approximate individual $E[I_t^a]$ as discussed in Section 3, these are all easy adjustments to the two algorithms presented.

For a more detailed discussion on accommodating a deterministic lead time in the context of inventory systems similar to the one discussed in this work we redirect the reader to Rossi et al. (2010), in particular to their Appendix A.1.

Appendix II: further details on our computational study

In this section we provide a table with details on the demand patterns used in our computational experience (Table 3) and a table of the 54 instances randomly selected in our test bed (Table 4).

Table 3. Demand patterns

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
STA	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
LCY1	0.54	0.72	0.96	1.22	1.54	1.86	2.2	2.52	2.82	3.06	3.24	3.32	3.32	3.24	3.06
LCY2	3.06	3.24	3.32	3.32	3.24	3.06	2.82	2.52	2.2	1.86	1.54	1.22	0.96	0.72	0.54
SIN1	2.42	2	1.58	1.4	1.58	2	2.42	2.6	2.42	2	1.58	1.4	1.58	2	2.42
SIN2	3.14	2	0.86	0.4	0.86	2	3.14	3.6	3.14	2	0.86	0.4	0.86	2	3.14
RAND	2.61	1.13	0.41	0.98	0.02	0.95	1.36	1.43	2.8	0.27	0.6	1.7	0.16	2.63	1.06
EMP1	0.01	0.25	0.76	2.33	1.34	2.44	2.23	1.24	1.40	1.81	0.77	1.46	1.1	0.46	0.53
EMP2	0.23	0.40	1.18	1.97	0.82	1.43	2.54	1.95	3.77	3.47	1.30	0.97	1.6	0.55	0.95
EMP3	0.22	0.58	1.32	0.72	0.73	0.99	0.37	0.91	1.02	0.57	0.82	1.59	0.59	2.41	2.67
EMP4	0.24	0.94	0.32	1.39	2.26	1.12	1.11	2.58	1.45	2.73	3.23	1.12	1.07	2.2	0.58

Table 4. The 54 instances randomly selected among the 270 in our test bed

Demand pattern	Order cost level	Penalty cost	Waste cost	Demand pattern	Order cost level	Penalty cost	Waste cost
EMP1		2.5	2	LCY1		5	2
EMP1		1	5	LCY1		2.5	5
EMP1		5	5	LCY1		5	10
EMP1		5	10	LCY1		5	10
EMP1		1	5	LCY2		2.5	5
EMP1		1	10	LCY2		2.5	10
EMP1		5	10	LCY2		1	10
EMP1		5	2	RAND		2.5	10
EMP1		5	5	RAND		1	10
EMP1		2.5	10	SIN1		5	2
EMP2		1	10	SIN1		2.5	10
EMP2		1	2	SIN1		5	5
EMP2		2.5	2	SIN1		1	10
EMP2		1	10	SIN1		1	10
EMP2		2.5	5	SIN2		5	2
EMP3		1	2	SIN2		2.5	2
EMP3		1	5	SIN2		1	10
EMP3		5	10	SIN2		5	5
EMP3		1	2	SIN2		1	10
EMP3		1	5	STA		1	2
EMP3		5	5	STA		2.5	2
EMP3		2.5	2	STA		2.5	2
EMP3		2.5	5	STA		5	2
EMP3		5	10	STA		5	5
EMP4		1	10	STA		2.5	5
EMP4		5	10	STA		2.5	10
LCY1		2.5	10	STA		2.5	10

References

- Alcoba, Alejandro G., Eligius M.T. Hendrix, Inmaculada Garcia, Gloria Ortega, Karin G.J. Pauls-Worm, and Rene Haijema. 2015. "On Computing Order Quantities for Perishable Inventory Control with Non-stationary Demand." In *Computational Science and Its Applications ICCSA 2015*, Vol. 9156 of *Lecture Notes in Computer Science* edited by Osvaldo Gervasi et al., 429–444. Springer, Cham.
- Askin, Ronald G. 1981. "A Procedure for Production Lot Sizing with Probabilistic Dynamic Demand." *A I I E Transactions* 13 (2): 132–137.
- Bakker, Monique, Jan Riezebos, and Ruud H. Teunter. 2012. "Review of inventory systems with deterioration since 2001." *European Journal of Operational Research* 221 (2): 275–284.
- Bollapragada, Srinivas, and Thomas E. Morton. 1999. "A Simple Heuristic for Computing Nonstationary (S, S) Policies." *Oper. Res.* 47 (4): 576–584.
- Bookbinder, James H., and Jin-Yan Tan. 1988. "Strategies for the Probabilistic Lot-Sizing Problem with Service-Level Constraints." *Management Science* 34 (9): 1096–1108.
- Caliskan-Demirag, Ozgun, Youhua F. Chen, and Yi Yang. 2012. "Ordering Policies for Periodic-Review Inventory Systems with Quantity-Dependent Fixed Costs." *Operations Research* 60 (4): 785–796.

- Chen, Xin, Zhan Pang, and Limeng Pan. 2014. "Coordinating Inventory Control and Pricing Strategies for Perishable Products." *Operations Research* 62 (2): 284–300.
- Federgruen, Awi, and Min Wang. 2015. "Inventory Models with Shelf-Age and Delay-Dependent Inventory Costs." *Operations Research* 63 (3): 701–715.
- Hendrix, Eligius M. T., Rene Haijema, Roberto Rossi, and Karin G. Pauls-Worm. 2012. "On Solving a Stochastic Programming Model for Perishable Inventory Control." In *Computational Science and Its Applications ICCSA 2012*, Vol. 7335 of *Lecture Notes in Computer Science* edited by Beniamino Murgante, et al., 45–56. Springer Heidelberg.
- Karaesmen, Itir Z., Alan Scheller-Wolf, and Borgia Deniz. 2011. "Managing Perishable and Aging Inventories: Review and Future Research Directions." In *Planning Production and Inventories in the Extended Enterprise*, Vol. 151 of *International Series in Operations Research & Management Science* edited by Karl G. Kempf, Pinar Keskinocak, and Reha Uzsoy, 393–436. Springer US.
- Li, Qing, and Peiwen Yu. 2014. "Multimodularity and Its Applications in Three Stochastic Dynamic Inventory Problems." *Manufacturing & Service Operations Management* 16 (3): 455–463.
- Minner, Stefan, and Sandra Transchel. 2010. "Periodic review inventory-control for perishable products under service-level constraints." *OR Spectrum* 32 (4): 979–996.
- Nahmias, Steven. 1982. "Perishable Inventory Theory: A Review." *Operations Research* 30 (4): 680–708.
- Pauls-Worm, Karin G.J., Eligius M.T. Hendrix, Alejandro G. Alcoba, and Rene Haijema. 2015. "Order quantities for perishable inventory control with non-stationary demand and a fill rate constraint." *International Journal of Production Economics* –.
- Pauls-Worm, Karin G. J., Eligius M. T. Hendrix, René Haijema, and Jack G. A. J. van der Vorst. 2014. "An MILP approximation for ordering perishable products with non-stationary demand and service level constraints." *International Journal of Production Economics* 157: 133–146.
- Rossi, Roberto. 2013. "Periodic Review for a Perishable Item under Non Stationary Stochastic Demand." In *Manufacturing Modelling, Management, and Control, Volume 7, Part 1*, edited by Bakhtadze Natalia, Jun., 2021–2026. International Federation of Automatic Control.
- Rossi, Roberto, Onur A. Kilic, and S. Armagan Tarim. 2015. "Piecewise linear approximations for the static-dynamic uncertainty strategy in stochastic lot-sizing." *Omega* 50: 126–140.
- Rossi, Roberto, S. Armagan Tarim, Brahim Hnich, and Steven Prestwich. 2010. "Computing the non-stationary replenishment cycle inventory policy under stochastic supplier lead-times." *International Journal of Production Economics* 127 (1): 180–189.
- Rossi, Roberto, S. Armagan Tarim, Steven Prestwich, and Brahim Hnich. 2014. "Piecewise linear lower and upper bounds for the standard normal first order loss function." *Applied Mathematics and Computation* 231: 489–502.
- Silver, Edward A. 1978. "Inventory control under a probabilistic time-varying, demand pattern." *A I I E Transactions* 10 (4): 371–379.
- Silver, Edward A., and Harlan C. Meal. 1973. "A heuristic for selecting lot size quantities for the case of a deterministic time-varying demand rate and discrete opportunities for replenishment." *Production and Inventory Management* 14 (2): 64–74.
- Tarim, S. Armagan, and Brian G. Kingsman. 2006. "Modelling and computing (Rn,Sn) policies for inventory systems with non-stationary stochastic demand." *European Journal of Operational Research* 174 (1): 581–599.
- Tunc, Huseyin, Onur A. Kilic, S. Armagan Tarim, and Burak Eksioglu. 2011. "The cost of using stationary inventory policies when demand is non-stationary." *Omega* 39 (4): 410–415.
- Wagner, Harvey M., and Thomson M. Whitin. 1958. "Dynamic Version of the Economic Lot Size Model." *Management Science* 5 (1): 89–96.
- Yang, Yi, Youhua F. Chen, and Yun Zhou. 2014. "Coordinating Inventory Control and Pricing Strategies Under Batch Ordering." *Operations Research* 62 (1): 25–37.