Computing the non-stationary replenishment cycle inventory policy under stochastic supplier lead-times

Citation for published version:
Rossi, R, Tarim, SA, Hnich, B & Prestwich, S 2010, 'Computing the non-stationary replenishment cycle inventory policy under stochastic supplier lead-times', *International Journal of Production Economics*, vol. 127, no. 1, pp. 180-189. https://doi.org/10.1016/j.ijpe.2010.05.011

Digital Object Identifier (DOI):
10.1016/j.ijpe.2010.05.011

Computing the Non-Stationary Replenishment Cycle Inventory Policy under Stochastic Supplier Lead-Times

Roberto Rossi [a,*] S. Armagan Tarim [b] Brahim Hnich [c] Steven Prestwich [d]

[a] *Logistics, Decision and Information Sciences, Wageningen UR, the Netherlands*

[b] *Operations Management Division, Nottingham University Business School, UK*

[c] *Faculty of Computer Science, Izmir University of Economics, Izmir, Turkey*

[d] *Cork Constraint Computation Centre, University College, Cork, Ireland*

Abstract

In this paper we address the general multi-period production/inventory problem with non-stationary stochastic demand and supplier lead time under service-level constraints. A replenishment cycle policy (R^n,S^n) is modeled, where R^n is the n-th replenishment cycle length and S^n is the respective order-up-to-level. We propose a Stochastic Constraint Programming approach for computing the optimal policy parameters. In order to do so, a dedicated global chance-constraint and the respective filtering algorithm that enforce the required service level are presented. Our numerical examples show that a stochastic supplier lead time significantly affects the structure of the optimal policy with respect to the case in which the lead time is assumed to be deterministic or absent.

Key words: inventory control; demand uncertainty; supplier lead time uncertainty; stochastic constraint programming; global chance-constraints

---

* **Corresponding author**. Roberto Rossi, Logistics, Decision and Information Sciences Group, Wageningen UR, Hollandseweg 1, 6706 KN, Wageningen, The Netherlands. Tel. +31 (0) 317 482321, Fax. +31 (0)317 485646.

Email addresses: `roberto.rossi@wur.nl` (Roberto Rossi), `armtar@yahoo.com` (S. Armagan Tarim), `brahim.hnich@ieu.edu.tr` (Brahim Hnich), `s.prestwich@4c.ucc.ie` (Steven Prestwich).

1 Introduction

An interesting class of production/inventory control problems is the one that considers the single location, single product case under stochastic demand. One of the well-known policies that can be adopted to control such a system is the "replenishment cycle policy", (R,S). Under the nonstationary demand assumption this policy takes the form (R^n,S^n), where R^n denotes the length of the nth replenishment cycle, and S^n the order-up-to-level value for the nth replenishment. This easy to implement inventory control policy yields at most $2N$ policy parameters fixed at the beginning of an N-period planning horizon. For a discussion on inventory control policies see Silver et al. [20]. The replenishment cycle policy provides an effective means of dampening the planning instability. Furthermore, it is particularly appealing when items are ordered from the same supplier or require resource sharing. In such a case all items in a coordinated group can be given the same replenishment period. Periodic review also allows a reasonable prediction of the level of the workload on the staff involved and is particularly suitable for advanced planning environments. For these reasons, as stated by Silver et al. [20], (R, S) is a popular inventory policy. Due to its combinatorial nature, the computation of (R^n,S^n) policy parameters is known to be a difficult problem to solve to optimality. An early approach proposed by Bookbinder and Tan [5] is based on a two-step heuristic method. Tarim and Kingsman [23,24] and Tempelmeier [27] propose a mathematical programming approach to compute policy parameters. Tarim and Smith [26] give a computationally efficient Constraint Programming formulation. An exact formulation and a solution method are presented in Rossi et al. [19].

All the above mentioned works assume either zero or a fixed (deterministic) supplier lead-time (i.e., replenishment lead-time). However, the lead-time uncertainty, which in various industries is an inherent part of the business environment, has a detrimental effect on inventory systems. For this reason, there is a vast inventory control literature analysing the impact of supplier lead-time uncertainty on the ordering policy (Whybark and Williams [29], Speh and Wagenheim [21], Nevison and Burstein [14]). A comprehensive discussion on stochastic supplier lead-time in continuous-time inventory systems is presented in Zipkin [30]. Kaplan [13] characterises the optimal policy for a dynamic inventory problem where the lead-time is a discrete random variable with known distribution and the demands in successive periods are assumed to form a stationary stochastic process. Since tracking all the outstanding orders through the use of Dynamic Programming requires a large multi-dimensional state vector, Kaplan assumes that orders do not cross in time and supplier lead time probabilities are independent of the size/number of outstanding orders (for details on order-crossover see Hayya et al. [9]).

The assumption that orders do not cross in time is valid for systems where the supplier production system has a single-server queue structure operating under a FIFO policy. Nevertheless, there are settings in which this assumption is not valid and orders cross in time. This has been recently investigated in Hayya et al. [8], Bashyam and Fu [3] and Riezebos [17]. As Riezebos underscores, the types of industries that have a higher probability of facing order crossovers are either located upstream in the supply chain, or use natural resources, or order strategic materials from multiple suppliers or from abroad. In a case study, he showed that the potential cost savings realized by taking order crossovers into account were in the order of 30%. Unfortunately, he remarks, modern ERP systems are not able to handle order crossovers effectively.

In a recent work, Babaï et al. [2] analyze a dynamic re-order point control policy for a single-stage, single-item inventory system with non-stationary demand and lead-time uncertainty. To the best of our knowledge, there is no complete or heuristic approach in the literature that addresses the computation of (R^n,S^n) policy parameters under stochastic supplier lead time and service level constraints. Computing optimal policy parameters under these assumptions is a hard problem from a computational point of view. We argue that incorporating both a non-stationary stochastic demand and a stochastic supplier lead time — without assuming that orders do not cross in time — in an optimization model is a relevant and novel contribution.

In this work, we propose a Stochastic Constraint Programming [28] model for computing optimal (R^n,S^n) policy parameters under service level constraints and stochastic supplier lead times. In Stochastic Constraint Programming, complex non-linear relations among decision and stochastic variables — such as the chance-constraints that enforce the required service level — can be effectively modeled by means of *global chance-constraints* [10]. Examples of global chance-constraints applied to inventory control problems can be found in [19,22]. Our model incorporates a dedicated global chance-constraint that enforces, for each replenishment cycle scheduled, the required non-stockout probability. The model is tested on a set of instances that are solved to optimality under a discrete stochastic supplier lead time with known distribution.

The paper is organized as follows. In Section 2 we provide the formal definition of the problem and we discuss the working assumptions. In Section 3 we provide a deterministic reformulation for the chance-constraints that enforce the required service level. In Section 4 we introduce Stochastic Constraint Programming and we discuss how it is possible to embed the deterministic reformulation of the chance-constraints within a global chance-constraint. This global chance-constraint is then enforced in the Stochastic Constraint Programming model for computing the optimal policy parameters. In Section 5 we present our computational experience on a set of instances. Finally, in Section 6, we draw conclusions.

2 Problem Definition

We consider the uncapacitated, single location, single product inventory problem with a finite planning horizon of N periods and a demand d_t for each period $t \in \{1, \ldots, N\}$, which is a random variable with probability density function $g_t(d_t)$. We assume that the demand occurs instantaneously at the beginning of each time period. The demand we consider is nonstationary, that is it can vary from period to period, and we also assume that demands in different periods are independent.

Following Eppen and Martin [6], an order placed in period t will be received after l_t periods, where l_t is a discrete random variable with probability mass function $f_t(\cdot)$. This means that an order placed in period t will be received after k periods with probability $f_t(k)$. We shall assume that there is a maximum lead-time L for which $\sum_{k=0}^{L} f_t(k) = 1$. Therefore the possible lead-time lengths are limited to $\Lambda = \{0, \ldots, L\}$ and the probability mass function is defined on the finite set Λ. Note that lead-times are mutually independent and each of them is also independent of the respective order quantity.

A fixed delivery cost a is incurred for each order. A linear holding cost h is incurred for each unit of product carried in stock from one period to the next. Without loss of generality, we will adopt the following assumption that concerns the accounting of inventory holding costs: we will charge an inventory holding cost at the end of each period based on the current inventory position, rather than the current inventory level. This will reflect the fact that interests are charged not only on the actual amount of items in stock, but also on outstanding orders. Doing so often makes sense since companies may assess holding cost on their total invested capital and not simply on items in stock. A further and detailed justification for this can be found in [11].

We assume that it is not possible to sell back excess items to the vendor at the end of a period and that negative orders are not allowed, so that if the actual stock exceeds the order-up-to-level for that review, this excess stock is carried forward and not returned to the supply source. However, such occurrences are regarded as rare events (see the discussion in [5,23]) and accordingly the cost of carrying this excess stock and its effect on the service levels of subsequent periods are ignored.

As a service level constraint we require that, with a probability of at least a given value α, at the end of each period the net inventory will be non-negative. Our aim is to minimize the expected total cost, which is composed of ordering costs and holding costs, over the N-period planning horizon, satisfying the service level constraints by fixing the future replenishment periods and the corresponding order-up-to-levels at the beginning of the planning horizon.

The actual sequence of actions is adopted from Kaplan [13]. At the beginning of a period, the inventory on hand after all the demands from previous periods have been realized is known. Since we are assuming complete backlogging, this quantity may be negative. Also known are orders placed in previous periods which have not been delivered yet. On the basis of this information, an ordering decision is made for the current period. All the deliveries that are to be made during a period are assumed to be made immediately after this ordering decision and hence are on hand at the beginning of the period. To summarize there are three successive events at the beginning of each period. First, stock on hand and outstanding orders are determined. Second, an ordering decision is made on the basis of this information. Third, all supplier deliveries for the current period, including possibly the most recent orders, are received.

3 Nonstationary Stochastic Lead-Time

Let us denote the inventory position (the total amount of stock on hand plus outstanding orders minus back-orders) at the end of period t as P_t. It directly follows that

$$P_t = I_t + \sum_{\{k|1\leq k\leq t, l_k+k>t\}} X_k, \tag{1}$$

where I_t is the inventory level (stock on hand minus back-orders) at the end of period t, X_k is the size of the replenishment order placed in period k, $X_k \geq 0$ (received in period $k+l_k$), and it is assumed that I_0 equals the initial inventory.

The general chance-constrained programming model for the problem described in Section 2 is given below. The reader is referred to Bookbinder and Tan [5] for the zero lead-time version of this problem.

$$\begin{aligned} \min \quad E\{TC\} = \int_{d_1} \dots \int_{d_N} \sum_{t=1}^{N} (a\delta_t + hP_t) \\ g_1(d_1)\dots g_N(d_N)\mathrm{d}(d_1)\dots \mathrm{d}(d_N) \end{aligned} \tag{2}$$

subject to,

$$\delta_t = \begin{cases} 1, & \text{if} \quad X_t > 0 \\ 0, & \text{otherwise} \end{cases} \qquad t = 1, ..., N \tag{3}$$

$$P_t = I_0 + \sum_{k=1}^{t}(X_k - d_k) \qquad t = 1, ..., N \tag{4}$$

$$\Pr\{P_t \geq \sum_{\{k|1\leq k\leq t, l_k>t-k\}} X_k\} \geq \alpha \qquad t = L+1, ..., N \tag{5}$$

$$P_t \in \mathbb{R}, \quad X_t \geq 0, \qquad t = 1, ..., N. \tag{6}$$

where we comply with the following notation:

$E\{.\}$: the expectation operator,
TC : total cost,
d_t : the demand in period t, a random variable with probability density function, $g_t(d_t)$,
a : the fixed ordering cost (incurred when an order is placed),
h : the proportional stock holding cost,
l_t : the lead-time length of the order placed in period t, a discrete random variable with a probability mass function $f_t(\cdot)$.
δ_t : a {0,1} variable that takes the value of 1 if a replenishment occurs in period t and 0 otherwise.

The objective function (Eq. 2) minimizes the expected total ordering and inventory holding cost. It should be noted that, by charging holding cost on the inventory position rather than on the inventory level, the objective function becomes particularly simple and it resembles the one employed when the lead time is zero. Eq. 3 states that if a replenishment occurs in period t — i.e. the order quantity X_t is greater than 0 — then the corresponding indicator variable δ_t must take value 1. Eq. 4 enforces the inventory conservation constraint for each period t, this constraint is expressed in terms of the inventory position P_t. Eq. 5 enforces the required service level in each period t, and it is also expressed in terms of the inventory position P_t. Finally Eq. 6 states that the inventory position in each period may either be zero or take any positive/negative value (i.e. full backorders) and that the order quantity is forced to be greater or equal to 0.

Note that depending on the probabilities assigned to each lead time length by the probability mass function, it may not be possible, in general, to provide the required service level for some initial periods. Nevertheless, by reasoning on a worst case scenario, it will always be possible to provide the required service level α starting from period $L+1$. Hence, the service level constraints are enforced in periods $L+1,\ldots,N$ (see Eq. 5).

Consider a review schedule, which has m reviews over the N period planning horizon with orders placed at $T_1, T_2, \ldots, T_m$, where $T_i < T_{i+1}$.In order to incorporate the "replenishment cycle policy" into this model, we express the whole model in terms of a new set of decision variables , R_{T_i}, $i = 1,\ldots,m$. Define,

$$P_t = R_{T_i} - \sum_{k=T_i}^{t} d_k, \quad T_i \le t < T_{i+1}, \quad i = 1,\ldots,m \tag{7}$$

where R_{T_i} ("order-up-to-position") can be interpreted as the inventory position up to which inventory should be raised after placing an order at the ith review period T_i. By doing so, order quantities X_t have to be decided only after the demands in the former periods have been realized. Under such a policy the orders X_t are all equal to zero except at replenishment periods $T_1, T_2, \ldots, T_m$.

The service level constraint has to be expressed as a relation between the order-up-to-positions such that the overall service level provided at the end of each period is at least α. In order to express this service level constraint we propose a scenario based approach over the discrete random variables l_t, $t = 1, \ldots, N$. In a scenario based approach [4,25], a scenario tree is generated which incorporates all possible realisations of discrete random variables into the model explicitly, yielding a fully deterministic model under the nonanticipativity constraints.

In our problem we can divide random variables into two sets: the random variables $\{l_t | t = 1, \ldots, N\}$, which represent lead-times, and the random variables $\{d_t | t = 1, \ldots, N\}$, which represent demands. We deal with each set in a separate fashion, by employing a scenario based approach for the l_t and a deterministic equivalent modeling approach for the d_t variables. This is possible since under a given scenario discrete random variables are treated as constants. The problem is then reduced to the general multi-period production/inventory problem with dynamic deterministic lead-times and stochastic demands. It should be noted that, although it has been assumed that the supplier lead-time is zero in Tarim and Kingsman [23], it is possible to extend their model for the non-zero lead-time situation without any loss of generality when the lead time is deterministic and remains constant for each order. In the Appendix we show how to model the situation in which the lead time is deterministic and dynamic (i.e. it may take a different deterministic value in each period). This more general situation corresponds to what is observed within any given scenario.

A scenario ω_t is a possible lead-time realization for all the orders placed up to period t in a given review schedule. We denote the probability of a scenario ω_t as $\Pr\{\omega_t\}$. Let $l_{T_i}(\omega_t)$ be the realized lead-time in scenario ω_t for the order placed in period T_i, where $i = 1, \ldots, m$. Finally, let Ω_t be the set of all the possible scenarios ω_t. Note that $\sum_{\Omega_t} \Pr\{\omega_t\} = 1$ for all $t = 1, \ldots, N$.

We define $T_{p(t)}$ as the latest review before period t in the planning horizon, for which we are sure that all the former orders, including the one placed in $T_{p(t)}$, have been delivered within period t. Under the assumption that the probability mass function $f_t(\cdot)$ is defined on a finite set Λ, the index $p(t)$ provides a bound for the scenario tree size. In fact if the possible lead-time lengths in Λ are $0, \ldots, L$, the earliest order that is delivered in period t with probability 1 under every possible scenario ω_t is the latest placed in the span $1, \ldots, t - L$. Therefore since each scenario ω_t identifies the orders that have been received before or in period t, it directly follows that the number of scenarios in the tree that is needed to compute the order-up-to-positions for periods $t - L, \ldots, t$ under any possible review schedule is at most 2^L, when we place $L + 1$ orders in periods $t - L, \ldots, t$, but it may be lower if fewer reviews are planned.

In order to clarify this, we shall provide a small numerical example. Consider a planning horizon of $N = 6$ periods. The probability mass function for the lead-time in each period $t = 1, \dots, 6$ is $f_t(\cdot) = \{0(1/3), 1(1/3), 2(1/3)\}$, therefore an order will arrive immediately with probability 1/3, after one period with probability 1/3, and after 2 periods with probability 1/3. It follows that in our example $L = 2$ and $f_t(\cdot)$ is defined on a finite set Λ that comprises 3 possible options. Let us now consider period $t = 5$. Clearly, $T_{p(t)} = 3$, in fact with probability 1.0 an order placed at period 3, as well as any other order placed at previous periods, is received by period 5. Under a review schedule that places an order in every period, there are $2^L = 4$ possible scenarios for the remaining orders that have been delivered by period 5:

- S_1, $\Pr\{S_1\} = (1/3 + 1/3)1/3$; both the orders placed at period 4 and 5 have been delivered by period 5.
- S_2, $\Pr\{S_2\} = (1/3 + 1/3)(1/3 + 1/3)$; the order placed at period 4 has been delivered by period 5, but not the one placed at period 5;
- S_3, $\Pr\{S_3\} = 1/3 \cdot 1/3$; the order placed at period 5 has been delivered by period 5, but not the one placed at period 4;
- S_4, $\Pr\{S_3\} = 1/3(1/3 + 1/3)$; the orders placed at period 4 and at period 5 have not been delivered by period 5;

It is easy to see that under any other possible review schedule the number of scenarios to be considered for the orders that have been delivered by period 5 is less or equal to $2^L = 4$. For instance, consider a review schedule in which orders are placed only in period 1, period 3, and period 5. In this case we only have 2 possible scenarios at period 5. As in the previous case, any order placed at period 3 or before will be received with probability 1.0 by period 5. No order is placed at period 4. The 2 scenarios for the remaining order are

- S_1, $\Pr\{S_1\} = 1/3$; the order placed at period 5 has been delivered by period 5;
- S_2, $\Pr\{S_1\} = 2/3$; the order placed at period 5 has not been been delivered by period 5.

The service level constraint at period t is always a relation over at most $L+1$ decision variables R_{T_i} that represent the order-up-to-positions of the replenishment cycles covering the span $t - L, \dots, t$. Let $p_\omega(t)$ be the value of $p(t)$ under a given scenario ω_t when a review schedule is considered. In order to satisfy the service level constraints in our original model, we require that the overall service level under all the possible scenarios for each set of at most $L + 1$ decision variables is at least α or equivalently,

$$\sum_{\omega_t \in \Omega_t} \Pr\{\omega_t\} \cdot G_S \left(R_{T_{p_\omega(t)}} + \sum_{\{i | i > p_\omega(t), l_{T_i}(\omega_t) \leq t - T_i\}} (R_{T_i} - R_{T_{i-1}}) \right) \geq \alpha, \quad t = L+1, \dots, N, \tag{8}$$

where $S = \sum_{k=T_{p_\omega(t)}}^{t} d_k - \sum_{\{i|i>p_\omega(t), l_{T_i}(\omega_t) \leq t-T_i\}} (d_{T_{i-1}} + \ldots + d_{T_i-1})$, and $G_S(.)$ is the cumulative distribution function of S. Further details on the derivation of Eq. 8 are provided in the Appendix.

As the reader may notice, the service level constraints (Eq. 8) are now fully deterministic constraints expressed only in terms of the order-up-to-positions, R_{T_i}. This makes it possible to replace throughout the rest of the model the P_t variables with their expected values $\tilde{P}_t$, as originally proposed in Bookbinder and Tan [5], since these are only affecting the objective function in which we are considering expected values.

We can now express the whole model in terms of the new set of decision variables R_t, $t = 1, \ldots, N$. If there is no replenishment scheduled for period t, that is if $\delta_t = 0$, then R_t must be equal to the expected closing-inventory-position in period $t-1$, that is $R_t = \tilde{P}_{t-1}$. If there is a review T_i in period t, R_t is simply the order-up-to-position, R_{T_i}, for this review. Therefore, the set of the desired order-up-to-positions, $\{R_{T_i} | i = 1, \ldots, m\}$, as required for the solution to the problem, comprises those values of R_t for which $\delta_t = 1$.

Hence, the complete deterministic equivalent model under the replenishment cycle policy can be expressed as

$$\min \quad E\{TC\} = \sum_{t=1}^{N} \left(a\delta_t + h\tilde{P}_t \right) \tag{9}$$

subject to,

$$\delta_t = 0 \Rightarrow R_t = \tilde{P}_{t-1} \qquad t = 1, \ldots, N \tag{10}$$

$$R_t \geq \tilde{P}_{t-1} \qquad t = 1, \ldots, N \tag{11}$$

$$R_t = \tilde{P}_t + \tilde{d}_t \qquad t = 1, \ldots, N \tag{12}$$

Eq. 8 (service level constraints),

$$R_t \geq 0, \quad \tilde{P}_t \geq 0, \quad \delta_t \in \{0, 1\} \qquad t = 1, \ldots, N, \tag{13}$$

where $\{T_1, \ldots, T_m\} = \{t \in \{1, \ldots, N\} | \delta_t = 1\}$.

The model neatly resembles the original stochastic programming formulation. The reader can easily notice that, while the objective function and the remaining constraints in the model are now deterministic and linear — thus they can be easily modeled by means of existing mathematical programming packages — Eq. 8 is deterministic but non-linear and it cannot be implemented in a straightforward manner by using existing solvers. For this reason, in the following section, we will introduce a Stochastic Constraint Programming formulation that we will employ to solve the above model.

4 A Stochastic Constraint Programming Approach

In this section, we aim to propose a Stochastic Constraint Programming approach for modeling and solving the model discussed in the previous section. Firstly, we introduce the key concepts in Constraint Programming and Stochastic Constraint Programming, the extension of Constraint Programming that deals with problems of decision making under uncertainty. Secondly, we introduce our Stochastic Constraint Programming model.

4.1 Constraint Reasoning

Constraint Programming (CP) [1] is a declarative programming paradigm in which relations between decision variables are stated in the form of constraints. Informally speaking, constraints specify the properties of a solution to be found. The constraints used in constraint programming are of various kinds: logic constraints (i.e. "x or y is true", where x and y are boolean decision variables), linear constraints, and *global constraints* [16]. A global constraint captures a relation among a non-fixed number of variables. One of the most well known global constraints is the **alldiff** constraint [15], that can be enforced on a certain set of decision variables in order to guarantee that no two variables are assigned the same value. With each constraint, CP associates a *filtering algorithm* able to remove provably infeasible or suboptimal values from the domains of the decision variables that are constrained and, therefore, to enforce some degree of *consistency* (see [18]). These filtering algorithms are repeatedly called until no more values are pruned. This process is called *constraint propagation*. In addition to constraints and filtering algorithms, constraint solvers also feature some sort of heuristic *search engine* (e.g. a backtracking algorithm). During the search, the constraint solver exploits filtering algorithms in order to proactively prune parts of the search space that cannot lead to a feasible or to an optimal solution.

Stochastic Constraint Programming (SCP) was first introduced in [28] in order to model combinatorial decision problems involving uncertainty and probability. According to Walsh, SCP combines together the best features of CP (i.e. global constraints, search heuristics, filtering strategies, etc.) and of Stochastic Programming [12] (i.e. stochastic variables, chance-constraints, etc.). In addition to decision variables, SCP features stochastic variables. Furthermore, in SCP it is possible to capture complex non-linear relations among decision and stochastic variables by means of *global chance-constraints* [19,10]. Similarly to global constraints, global chance-constraints incorporate efficient strategies for performing logical inference on these relations during the search in order to enforce some degree of consistency through constraint propagation.

In what follows we will introduce an SCP model for computing (R^n,S^n) policy parameters under non-stationary stochastic demand, lead time, and service level constraints. In order to capture the service level constraints, a dedicated global chance-constraint and the respective propagation logic are introduced and incorporated in the SCP model.

4.2 A Stochastic Constraint Programming Model

We now present an SCP formulation for computing (R^n,S^n) policy parameters under stochastic lead times. Results from Section 3 will be employed in the SCP formulation. More specifically, in order to model the service level constraint (Eq. 8), a new global chance-constraint, `serviceLevel`$(\cdot)$, will be defined. Such a constraint is needed to dynamically compute the correct expected closing-inventory-positions $\{\tilde{P}_t | t = 1, \ldots, N\}$ on the basis of the current replenishment plan, that is $\{\delta_t | t = 1, \ldots, N\}$ assignments.

The SCP model that incorporates our dedicated global chance-constraint is therefore

$$\min \quad E\{TC\} = \sum_{t=1}^{N} \left(a \cdot \delta_t + h \cdot \tilde{P}_t \right) \tag{14}$$

subject to,

$$\delta_t = 0 \Rightarrow \tilde{P}_t + \tilde{d}_t - \tilde{P}_{t-1} = 0 \qquad t = 1, \ldots, N \tag{15}$$

$$\tilde{P}_t + \tilde{d}_t - \tilde{P}_{t-1} \geq 0 \qquad t = 1, \ldots, N \tag{16}$$

$$\begin{aligned} &serviceLevel(\delta_1, \ldots, \delta_N, \tilde{P}_1, \ldots, \tilde{P}_N, \\ &\qquad g_1(d_1), \ldots, g_N(d_N), f(\cdot), \alpha) \end{aligned} \tag{17}$$

$$\tilde{P}_t \geq 0, \quad \delta_t \in \{0, 1\} \qquad t = 1, \ldots, N. \tag{18}$$

It should be noted that the domain of each $\tilde{P}_t$ variable — as in the zero lead time case (see Tarim and Smith [26]) — is limited. In fact, since the period demand variance is additive, the uncertainty can only increase in the length of a replenishment cycle. Therefore the longer a cycle is, the higher are the inventory levels that are required to achieve a certain service level. It directly follows that a single replenishment covering the whole planning horizon will provide upper bounds for the expected period closing-inventory-positions throughout the horizon.

We now describe the signature of the new constraint we have introduced. `serviceLevel`$(\cdot)$ describes a relation between all the decision variables in the model. It also accepts as parameters the distribution of the demand in each period t, $g(d_t)$; the probability mass function of the lead time $f(\cdot)$, which, without loss of generality, is here assumed to be the same for all the periods;

and the required service level α.

A high level pseudo-code for the propagation logic of `serviceLevel`$(\cdot)$ is presented in Algorithm 1. Note that to keep the description of the algorithm simple we assume here a stochastic lead time l with probability mass function $f(l)$ in every period. The maximum lead time length is L.

Algorithm 1: `propagate`

input : $\delta_1, \ldots, \delta_N$, $\tilde{P}_1, \ldots, \tilde{P}_N$, α, $d_1, \ldots, d_N$, l, L, N

```
1  begin
2      cycles ← {};
3      pointer ← 1;
4      periods ← 0;
5      for each period i in 2, ..., N do
6          if δ_i is not assigned then
7              cycles ← {};
8              periods ← 0;
9              pointer = −1;
10         else if δ_i is assigned to 1 then
11             if pointer ≠ −1 then
12                 cycle ← a replenishment cycle over {pointer, ..., i − 1};
13                 add cycle to cycles;
14             if periods ≥ L then
15                 checkBuffers();
16             pointer ← i;
17             periods ← periods + 1;
18         else
19             periods ← periods + 1;
20     if pointer ≠ −1 then
21         cycle ← a replenishment cycle over {pointer, ..., N};
22         add cycle to cycles;
23     if periods ≥ L then
24         checkBuffers();
25 end
```

In order to propagate this constraint, we consider every set of consecutive replenishment cycles covering at least $L+1$ periods (that is the one of interest plus L former periods) and having the smallest possible cardinality in terms of replenishment cycle number (Algorithm 1, line 5). Obviously, to identify such a group of cycles, we have to wait until, during the search, a subset of

Procedure `checkBuffers`

```
1  begin
2      cycle ← the last element in cycles, a replenishment cycle over {i, ..., j};
3      if no decision variable P̃_i, ..., P̃_j is assigned then
4          return;
5      counter ← 1;
6      for each period t covered by cycle do
7          formerCycles ← cycles;
8          remove cycle from formerCycles;
9          coveredPeriods ← the number of periods covered by cycles in
           formerCycles;
10         head ← first element in formerCycles;
11         headLength ← periods covered by head;
12         if counter < L then
13             while coveredPeriods − headLength + counter ≥ L do
14                 remove head from formerCycles;
15                 head ← first element in formerCycles;
16                 headLength ← periods covered by head;
17         else
18             formerCycles ← {};
19         condition ← true;
20         for each cycle c in formerCycles do
21             let {m, ..., n} be the periods covered by c;
22             if no decision variable P̃_m, ..., P̃_n is assigned then
23                 condition ← false;
24         if condition then
25             if Eq. 8 for period t in cycle and former replenishment
               cycles in formerCycles is not satisfied then
26                 backtrack() ;
27         counter ← counter + 1;
28 end
```

consecutive δ_t variables is assigned (Algorithm 1, line 10). Then, in order to verify if the service level constraint is satisfied for the last period in this group, we check that for each replenishment cycle in the group identified at least one decision variable $\tilde{P}_t$ is assigned (Procedure `checkBuffers`, line 3 and line 22). If this is the case the partial policy for the span is completely defined and, by recalling that $R_t = \tilde{P}_t + \tilde{d}_t$, its feasibility can be checked by using the condition in Eq. 8 (Procedure `checkBuffers`, line 25). If the condition is not

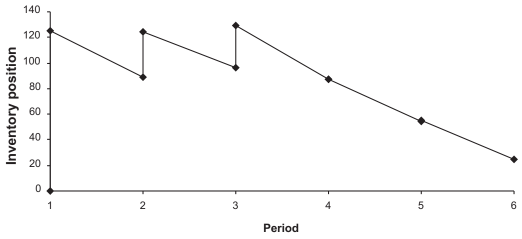

Fig. 1. Optimal policy under stochastic lead time, $f_t(k) = \{0.3, 0.2, 0.5\}$.

$E\{TC\}$: 356					
Period (t)	1	2	3	4	5
$\tilde{d}_t$	36	28	42	33	30
R_t	125	124	129	87	55
δ_t	1	1	1	1	1
Shortage probability	–	–	5%	5%	5%

Table 1
Optimal solution.

satisfied we backtrack (Procedure `checkBuffers`, line 26). Notice that such a condition involves for each period only a subset of all the decision variables in the model, which means that our constraint is able to detect infeasible partial assignments, i.e. *nogoods* [18].

Finally, it should be emphasized that, during the search, any CP solver will be able to exploit constraint propagation and detect infeasible or suboptimal assignments with respect to the other constraints in the model. Furthermore, suboptimal solutions may be pruned by using dedicated *cost-based filtering* methods [7,22].

4.2.1 An example

We assume an initial null inventory level and a normally distributed demand with a coefficient of variation $\sigma_t/\tilde{d}_t = 0.3$ for each period $t \in \{1, \ldots, 5\}$. The expected values for the demand in each period are: {36, 28, 42, 33, 30}. The other parameters are $a = 1$, $h = 1$, $\alpha = 0.95$. We consider for every period t in the planning horizon the following lead time probability mass function $f_t(k) = \{0.3(0), 0.2(1), 0.5(2)\}$, which means that we receive an order placed in period t after $\{0, \ldots, 2\}$ periods with the given probability (0 periods: 30%; 1 period: 20%; 2 periods: 50%). It is obvious that in this case we will always receive the order at most after 2 periods. In Table 1 (Fig. 1) we show the optimal solution found by the SCP model. We now want to show that the order-up-to-positions — computed in this example by using Eq. 8 — satisfy every service level constraint in the model. We assume that for the first 2 periods no service level constraint is enforced, since it is not possible to fully control the inventory in the first 2 periods. Therefore we enforce the required service level on periods 3, 4 and 5, that is Eq. 8 for $t = 3, \ldots, N$. Let us verify that the given order-up-to levels satisfy this condition for each of these three periods. Since we know the probability mass function $f_t(\cdot)$ for each period t in the planning horizon we can easily compute the probability $\Pr(\omega_t)$ for each scenario $\omega_t \in \Omega_t$. We have four of these scenarios for each period $t \in \{3, \ldots, N\}$, since we are placing an order in every period:

- S_1, $\Pr\{S_1\} = 0.15 = (0.3+0.2)0.3$; in this scenario at period t all the orders placed are received. That is the order placed in period $t-1$ is received immediately (probability 0.3), or after one period (probability 0.2), while the order placed in period t is received immediately (probability 0.3)
- S_2, $\Pr\{S_2\} = 0.35 = (0.3+0.2)(0.2+0.5)$; in this scenario at period t we do not receive the last order placed in period t. That is the order placed in period $t-1$ is received immediately (probability 0.3), or after one period (probability 0.2), while the order placed in period t is not received immediately, therefore it is received after one period (probability 0.2), or after two periods (probability 0.5)
- S_3, $\Pr\{S_3\} = 0.35 = 0.5(0.2+0.5)$; in this scenario at period t we don't receive the last two orders placed in periods t and $t-1$. That is the order placed in period $t-1$ is received after two periods (probability 0.5), and the order placed in period t is not received immediately, therefore it is received after one period (probability 0.2), or after two periods (probability 0.5)
- S_4, $\Pr\{S_4\} = 0.15 = 0.5 \cdot 0.3$; in this scenario at period t we don't receive the order placed in period $t-1$ and we observe order-crossover. That is the order placed in period $t-1$ is received after two periods (probability 0.5), and the order placed in period t is received immediately (probability 0.3)

In the described scenarios every possible configuration is considered. We do this without any loss of generality. In fact if some of the configurations are unrealistic (for instance if we assume that order-crossover may not take place) we just need to set the probability of the respective scenario to zero. Now it is possible to write Eq. 8 for each period $t \in \{3, \ldots, N\}$. Consider period 3:

$$\begin{aligned}\Pr\{S_1\} \cdot G\left(\frac{129-42}{0.3\sqrt{42^2}}\right) + \Pr\{S_2\} \cdot G\left(\frac{124-(28+42)}{0.3\sqrt{28^2+42^2}}\right) &+ \\ \Pr\{S_3\} \cdot G\left(\frac{125-(36+28+42)}{0.3\sqrt{36^2+28^2+42^2}}\right) &+ \\ \Pr\{S_4\} \cdot G\left(\frac{125+(129-124)-(36+42)}{0.3\sqrt{36^2+42^2}}\right) = 94.60\% \cong 95\% &\end{aligned} \tag{19}$$

where $G(\cdot)$ is the standard normal distribution function. This means that the combined effect of order delivery delays in our policy, when all the possible scenarios are taken into account, gives a no stock-out probability of about 95% for period 3. A similar reasoning can be employed to verify that the given solution satisfies the required service level also for period $t \in \{4, 5\}$.

$E\{TC\}$: 211 (lower bound)					
Period (t)	1	2	3	4	5
$\tilde{d}_t$	36	28	42	33	30
R_t	–	124	100	87	–
δ_t	–	1	1	1	–
Shortage probability				6%	

Table 2
A partial assignment and the respective shortage probability in period 4. The dashes, "-", are used to denote decision variables that have not been assigned yet.

Period (t)	1	2	3	4	5	6	7	8
$\tilde{d}_t$	15	18	13	33	30	18	23	15

Table 3
Forecasts of period demands.

The reader may notice that, since we are placing an order in every period and since the lead time is at most of two periods, the service level in any given period is only influenced by the replenishment in such a period and by the last two replenishments. For instance, the service level in period 4 is only influenced by the order-up-to-position in periods 3 and 2. Let us consider the partial assignment in Table 2. The shortage probability in period 4 is greater than the required 5% therefore this partial assignment constitutes a *nogood*. As soon as our global chance-constraint detects this partial assignment during the search, it will immediately trigger a backtrack and it will prevent the CP solver from exploring any assignment that extends such a partial assignment.

5 Computational Experience

In this section we solve to optimality an 8-period inventory problem under stochastic demand and lead time. Different lead time configurations are considered. The stochastic, deterministic and zero lead time cases are compared. As in the previous example we assume an initial null inventory level and a normally distributed demand with a coefficient of variation $\sigma_t/\tilde{d}_t = 0.3$ for each period $t \in \{1, \dots, 8\}$. The expected value $\tilde{d}_t$ for the demand in each period $t = 1, \dots, N$ are listed in Table 3. The other parameters are $a = 30$, $h = 1$, $\alpha = 0.95$. Initially we consider the problem under stochastic demand and no lead time, an efficient CP approach to find policy parameters in this case was presented in [26,22]. Obviously our approach is general and can provide solutions for this case as well, although less efficiently. The optimal solution for the instance considered is presented in Fig. 2, details about the optimal policy are reported in Table 4. We observe 5 replenishment cycles, policy parameters are: cycle lengths= $[1, 2, 1, 2, 2]$ and order-up-to-positions= $[72, 42, 49, 65, 52]$. The shortage probability is at most 5%, therefore the service level is met in every period. The $E\{TC\}$ is 303.

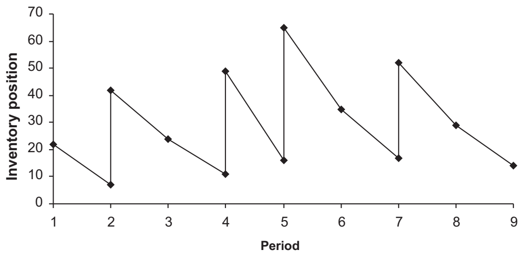

Fig. 2. Optimal policy under no lead time.

$E\{TC\}$: 303								
Period (t)	1	2	3	4	5	6	7	8
R_t	22	42	24	49	65	35	52	29
δ_t	1	1	0	1	1	0	1	0
Shortage probability	5%	0%	5%	5%	0%	5%	0%	5%

Table 4
Optimal policy under no lead time.

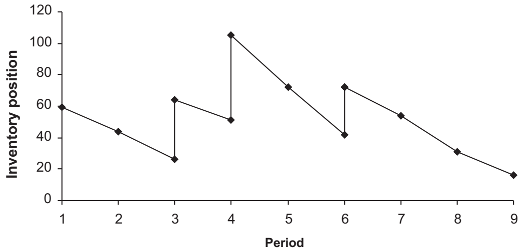

Fig. 3. Optimal policy under deterministic one period lead time.

$E\{TC\}$: 456								
Period (t)	1	2	3	4	5	6	7	8
R_t	59	44	64	105	72	72	54	31
δ_t	1	0	1	1	0	1	0	0
Shortage probability	–	0%	5%	5%	0%	5%	0%	5%

Table 5
Optimal policy under deterministic one period lead time, notice that the service level in the first period can obviously not be controlled.

We now consider the same instance, but with a deterministic lead time of one period. The optimal solution is presented in Fig. 3, details about the optimal policy are reported in Table 5. We observe now only 4 replenishment cycles, policy parameters are: cycle lengths= $[2, 1, 2, 3]$ and order-up-to-positions= $[59, 64, 105, 72]$. Again the shortage probability is at most 5% in every period, which means that the service level constraint is met. The $E\{TC\}$ is 456. Therefore we observe now an expected total cost that is 50.5% higher than the zero lead time case. The replenishment plan is significantly affected by the lead time both in term of replenishment cycle lengths and order-up-to-positions.

When a deterministic lead time of two periods is considered, as the reader may

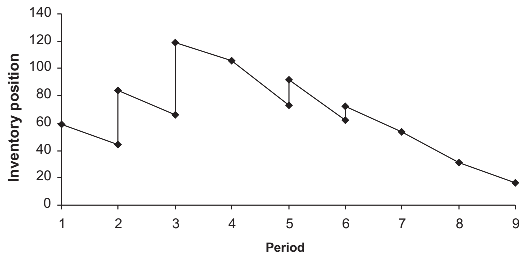

Fig. 4. Optimal policy under deterministic two periods lead time.

$E\{TC\}$: 602								
Period (t)	1	2	3	4	5	6	7	8
R_t	59	84	119	106	92	72	54	31
δ_t	1	1	1	0	1	1	0	0
Shortage probability	–	–	5%	5%	0%	5%	5%	5%

Table 6
Optimal policy under deterministic two periods lead time.

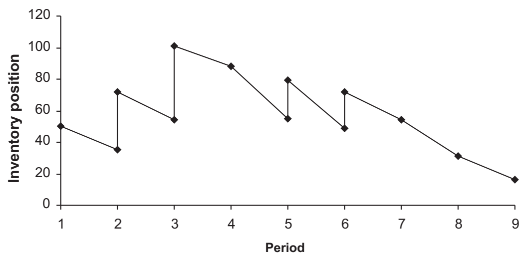

Fig. 5. Optimal policy under stochastic lead time, $f_t(k) = \{0.2(0), 0.6(1), 0.2(2)\}$.

expect, we observe again higher costs and a different replenishment policy. The optimal solution is presented in Fig. 4, details about the optimal policy are reported in Table 6. The number of replenishment cycles is now again 5, policy parameters are: cycle lengths= $[1, 1, 2, 1, 3]$ and order-up-to-positions= $[59, 84, 119, 92, 72]$. The service level constraint is met in every period. The $E\{TC\}$ is 602. This means that we observe a cost 98.6% and 32.0% higher than respectively the zero lead time case and the one period lead time case. The replenishment plan is again completely modified as a consequence of the lead time length.

We now concentrate on two instances where a stochastic lead time is considered and we compare results with the former cases. Firstly we analyze a stochastic lead time with probability mass function $f_t(k) = \{0.2(0), 0.6(1), 0.2(2)\}$. That is an order is received immediately with probability 0.2, after one period with probability 0.6, and after two periods with probability 0.2. The optimal solution is presented in Fig. 5, details about the optimal policy are reported in Table 7. The number of replenishment cycles is again 5 as in the two period lead time case, policy parameters are: cycle lengths= $[1, 1, 2, 1, 3]$ and order-up-to-positions= $[50, 72, 101, 79, 72]$. Therefore we see that the number and the length of replenishment cycles does not change from the deterministic

$E\{TC\}$: 532								
Period (t)	1	2	3	4	5	6	7	8
R_t	50	72	101	88	79	72	54	31
δ_t	1	1	1	0	1	1	0	0
Shortage probability	–	–	5%	5%	3%	5%	5%	5%

Table 7
Optimal policy under stochastic lead time, $f_t(k) = \{0.2(0), 0.6(1), 0.2(2)\}$, in periods $\{1, 2\}$ the inventory cannot be controlled.

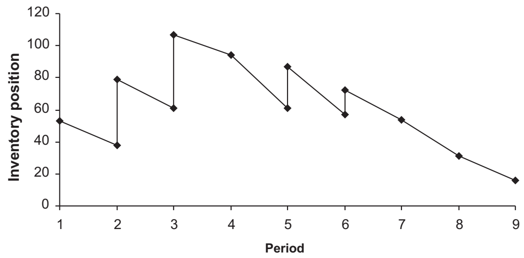

Fig. 6. Optimal policy under stochastic lead time, $f_i(t) = \{0.5(0), 0.0(1), 0.5(2)\}$.

$E\{TC\}$: 562								
Period (t)	1	2	3	4	5	6	7	8
R_t	53	79	107	94	87	72	54	31
δ_t	1	1	1	0	1	1	0	0
Shortage probability	–	–	5%	5%	0%	5%	5%	5%

Table 8
Optimal policy under stochastic lead time, $f_i(t) = \{0.5(0), 0.0(1), 0.5(2)\}$.

two period lead time case, although we observe lower order-up-to-positions as we may expect since the lead time is in average one period therefore lower than in the former case. Also the cost reflects this, in fact it is 11.6% lower than in the two period deterministic lead time case. It should be noted that the uncertainty of the lead time plays a significant role, in fact although the average lead time is one period, the structure of the policy resembles much more the one under a two period deterministic lead time than the one under a deterministic one period lead time. Moreover the expected total cost is 16.6% higher than in this latter case.

We finally consider a different probability mass function for the lead time: $f_t(k) = \{0.5(0), 0.0(1), 0.5(2)\}$, which means that we maintain the same average lead time of one period, but we increase its variance. The optimal solution is presented in Fig. 6, details about the optimal policy are reported in Table 8. The number of replenishment cycles is still 5, policy parameters are: cycle lengths= $[1, 1, 2, 1, 3]$ and order-up-to-positions= $[50, 72, 101, 79, 72]$. Although the average lead time is still one period, order-up-to-positions are slightly higher than in the former case where the variance of the lead time was lower. Also the cost reflects this, in fact it is 5.6% higher than in the former case, but still lower than the expected total cost of the two period deterministic lead time case.

To summarize, in our experiments we saw that supplier lead time uncertainty may significantly affect the structure of the optimal (R^n,S^n) policy. Computing optimal policy parameters constitutes a hard computational and theoretical challenge. Under different degrees of lead time uncertainty, when other input parameters for the problem remain fixed, order-up-to-positions and reorder points in the optimal policy change significantly. Realizing what the optimal decisions are for certain input parameters is a counterintuitive task. Our approach provides a systematic way to compute these optimal policy parameters.

6 Conclusions

A novel approach for computing replenishment cycle policy parameters under non-stationary stochastic demand, stochastic lead time and service level constraints has been presented. The approach is based on SCP and it employs a dedicated global chance-constraint in order to enforce the required service level in each period. The assumptions under which we developed our approach for the stochastic lead time case proved to be less restrictive than those commonly adopted in the literature for complete methods. In particular we faced the problem of order-crossover, which is a very active research topic. Our approach merged well known concepts such as deterministic equivalent modeling of chance-constraints and scenario based modeling. Our computational experience showed that a stochastic supplier lead time may significantly impact the structure and the cost of the optimal replenishment cycle policy with respect to the case in which the lead time is deterministic or absent. In our future research, we aim to develop dedicated cost-based filtering algorithms able to significantly speed up the search for the optimal policy parameters.

7 Appendix

In this Appendix we discuss the main steps required to derive the deterministic equivalent non-linear formulation of the service level constraints (Eq. 8).

To begin, we discuss how to obtain a deterministic equivalent formulation for the chance-constraints that enforce the required service level when the lead time in each period varies and assumes a given deterministic value. The same reasoning is then easily generalized to the case in which the lead time is stochastic and assumes a different distribution from period to period.

When a dynamic deterministic lead time $L_t \geq 0$ is considered in each period

$t = 1, \ldots, N$, an order placed in period t will be received only at period $t + L_t$, that is

$$I_t = I_0 + \sum_{\{k|k \geq 1, L_k + k \leq t\}} X_k - \sum_{k=1}^{t} d_k \qquad t = 1, ..., N. \tag{20}$$

Let us recall that the inventory position, P_t, represents the total amount of inventory on-hand plus outstanding orders minus backorders at the end of period t. It directly follows that

$$P_t = I_t + \sum_{\{k|1 \leq k \leq t, L_k + k > t\}} X_k, \tag{21}$$

where we assume $P_0 = I_0$. It is easy, then, to reformulate the model using the inventory position.

Next, we modify the general stochastic programming formulation in order to incorporate the "replenishment cycle policy". Consider a review schedule, which has m reviews over the N period planning horizon with orders placed at $\{T_1, T_2, \ldots, T_m\}$, where $T_i > T_{i-1}$, $T_m \leq N - L_{T_m}$. For convenience, T_1 is defined as the start of the planning horizon and $T_{m+1} = N + 1$ as the period immediately after the end of the planning horizon. [1] The associated inventory reviews will take place at the beginning of periods T_i, $i = 1, \ldots, m$. In the replenishment cycle policy considered here, clearly the orders X_t are all equal to zero except at replenishment periods $T_1, T_2, \ldots, T_m$. The inventory level I_t carried from period t to period $t + 1$ is the opening inventory plus any orders that have arrived up to and including period t less the total demand to date. Hence, the inventory balance equation becomes,

$$I_t = I_0 + \sum_{\{i|L_{T_i} + T_i \leq t\}} X_{T_i} - \sum_{k=1}^{t} d_k, \quad t = 1, \ldots, N. \tag{22}$$

Define $T_{p(t)}$ as the latest review before period t in the planning horizon, for which all the former orders, including the one placed in $T_{p(t)}$, are delivered within period t, therefore

$$p(t) = \max \left\{ i | \forall j \in \{1, \ldots, i\}, T_j + L_{T_j} \leq t, \quad i = 1, \ldots, m \right\}, \tag{23}$$

for all $t = 1, \ldots, N$. The inventory level I_t at the end of period t (Eq. 22) can be expressed as

$$I_t = I_0 + \sum_{i=1}^{p(t)} X_{T_i} + \sum_{\{i|i > p(t), L_{T_i} + T_i \leq t\}} X_{T_i} - \sum_{k=1}^{t} d_k, \quad t = 1, \ldots, N. \tag{24}$$

[1] The review schedule may be generalized to consider the case where $T_1 > 1$, if the opening inventory I_0 is sufficient to cover the immediate needs at the start of the planning horizon.

We now want to reformulate the constraints of the chance-constrained model in terms of a new set of decision variables R_{T_i}, $i = 1, \ldots, m$. Define,

$$P_t = R_{T_i} - \sum_{k=T_i}^{t} d_k, \quad T_i \leq t < T_{i+1}, \quad i = 1, \ldots, m \tag{25}$$

where R_{T_i} can be interpreted as the inventory position up to which inventory should be raised after placing an order at the ith review period T_i. We can now express the whole model in terms of these new decision variables R_{T_i}. The new problem is to determine the number of reviews, m, the T_i, and the associated R_{T_i} for $i = 1, \ldots, m$.

Let us now express Eq. 24 using R_{T_i} as decision variables

$$I_t = R_{T_{p(t)}} + \sum_{\{i|i>p(t), L_{T_i}+T_i \leq t\}} \left(R_{T_i} - R_{T_{i-1}} + d_{T_{i-1}} + \ldots + d_{T_i - 1} \right) - \sum_{k=T_{p(t)}}^{t} d_k, \quad t = 1, \ldots, N. \tag{26}$$

As mentioned earlier, α is the desired minimum probability that the net inventory level in any time period is non-negative. Depending on the values assigned to L_t it is obviously not possible to provide the required service level for some initial periods. In general, we provide the required service level α starting from the period t, for which the value $t + L_t$ is minimum. Let M be this period. Notice that, it will never be optimal to place any order in a period t such that $t + L_t > N$, since such an order will not be received within the given planning horizon.

By substituting I_t with the right hand term in Eq. 26 we obtain

$$G_S \left(R_{T_{p(t)}} + \sum_{\{i|i>p(t), L_{T_i}+T_i \leq t\}} (R_{T_i} - R_{T_{i-1}}) \right) \geq \alpha, \quad t = M, \ldots, N. \tag{27}$$

where $S = \sum_{k=T_{p(t)}}^{t} d_k - \sum_{\{i|i>p(t), L_{T_i}+T_i \leq t\}} (d_{T_{i-1}} + \ldots + d_{T_i - 1})$, and $G_S(.)$ is the cumulative distribution function of S.

The service level constraints are now deterministic and they are expressed only in terms of the order-up-to-positions. Eq. 27 can be directly employed in order to obtain Eq. 8, under the original assumption that the lead time in each period $t \in \{1, \ldots, N\}$ is a discrete random variable l_t.

References

[1] K. Apt. *Principles of Constraint Programming.* Cambridge University Press, Cambridge, UK, 2003.

[2] M. Z. Babaï, A. Syntetos, Y. Z. Dallery, and K. Nikolopoulos. Dynamic reorder point inventory control with lead-time uncertainty: analysis and empirical investigation. *International Journal of Production Research*, 47(9):2461–2483, 2009.

[3] S. Bashyam and M.C. Fu. Optimization of (s,S) inventory systems with random lead times and a service level constraint. *Management Science*, 44(12):243–256, 1998.

[4] J. R. Birge and F. Louveaux. *Introduction to Stochastic Programming.* Springer Verlag, New York, 1997.

[5] J. H. Bookbinder and J. Y. Tan. Strategies for the probabilistic lot-sizing problem with service-level constraints. *Management Science*, 34(9):1096–1108, 1988.

[6] G. D. Eppen and R. K. Martin. Determining safety stock in the presence of stochastic lead time and demand. *Management Science*, 34(11):1380–1390, 1988.

[7] F. Focacci, A. Lodi, and M. Milano. Cost-based domain filtering. In *Proceedings of the 5th International Conference on the Principles and Practice of Constraint Programming*, pages 189–203. Springer Verlag, 1999. Lecture Notes in Computer Science No. 1713.

[8] J. C. Hayya, U. Bagchi, J. G. Kim, and D. Sun. On static stochastic order crossover. *International Journal of Production Economics*, 114(1):404–413, July 2008.

[9] J. C. Hayya, S. H. Xu, R. V. Ramasesh, and X. X. He. Order crossover in inventory systems. *Stochastic Models*, 11(2):279–309, 1995.

[10] B. Hnich, R. Rossi, S. A. Tarim, and S. D. Prestwich. Synthesizing filtering algorithms for global chance-constraints. In *Principles and Practice of Constraint Programming, CP 2009, Proceedings*, volume 5732 of *Lecture Notes in Computer Science*, pages xxx–xxx. Springer, 2009. Forthcoming.

[11] J. A. Hunt. Balancing accuracy and simplicity in determining reorder points. *Management Science*, 12(4):B94–B103, 1965.

[12] P. Kall and S. W. Wallace. *Stochastic Programming.* John Wiley & Sons, 1994.

[13] R. S. Kaplan. A dynamic inventory model with stochastic lead times. *Management Science*, 16(7):491–507, 1970.

[14] C. Nevison and M. Burstein. The dynamic lot-size model with stochastic lead-times. *Management Science*, 30(1):100–109, 1984.

[15] J.-C. Régin. A filtering algorithm for constraints of difference in csps. In *Proceedings of the twelfth national conference on Artificial intelligence (vol. 1), Seattle, Washigton*, pages 362–367. American Association for Artificial Intelligence, 1994.

[16] J.-C Régin. *Global Constraints and Filtering Algorithms.* in Constraints and Integer Programming Combined, Kluwer, M. Milano editor, 2003.

[17] J. Riezebos. Inventory order crossovers. *International Journal of Production Economics*, 104(2):666–675, December 2006.

[18] F. Rossi, P. van Beek, and T. Walsh. *Handbook of Constraint Programming (Foundations of Artificial Intelligence).* Elsevier Science Inc., New York, NY, USA, 2006.

[19] R. Rossi, S. A. Tarim, B. Hnich, and S. Prestwich. A global chance-constraint for stochastic inventory systems under service level constraints. *Constraints*, 13(4):490–517, 2008.

[20] E. A. Silver, D. F. Pyke, and R. Peterson. *Inventory Management and Production Planning and Scheduling.* John-Wiley and Sons, New York, 1998.

[21] T. W. Speh and G. Wagenheim. Demand and lead-time uncertainty: The impacts of physical distribution performance and management. *Journal of Business Logistics*, 1(1):95–113, 1978.

[22] S. A. Tarim, B. Hnich, R. Rossi, and S. Prestwich. Cost-based filtering techniques for stochastic inventory control under service level constraints. *Constraints*, 14(2):137–176, 2009.

[23] S. A. Tarim and B. G. Kingsman. The stochastic dynamic production/inventory lot-sizing problem with service-level constraints. *International Journal of Production Economics*, 88(1):105–119, 2004.

[24] S. A. Tarim and B. G. Kingsman. Modelling and Computing (R^n,S^n) Policies for Inventory Systems with Non-Stationary Stochastic Demand. *European Journal of Operational Research*, 174(1):581–599, 2006.

[25] S. A. Tarim, S. Manandhar, and T. Walsh. Stochastic constraint programming: A scenario-based approach. *Constraints*, 11(1):53–80, 2006.

[26] S. A. Tarim and B. Smith. Constraint Programming for Computing Non-Stationary (R,S) Inventory Policies. *European Journal of Operational Research*, 189(3):1004–1021, 2008.

[27] H. Tempelmeier. On the stochastic uncapacitated dynamic single-item lotsizing problem with service level constraints. *European Journal of Operational Research*, 181(1):184–194, 2007.

[28] T. Walsh. Stochastic constraint programming. In *Proceedings of the 15th Eureopean Conference on Artificial Intelligence, ECAI'2002, Lyon, France, July 2002*, pages 111–115, 2002.

[29] D. C. Whybark and J. G. Williams. Material requirements planning under uncertainty. *Decision Science*, 7(4):595–606, 1976.

[30] P. Zipkin. Stochastic leadtimes in continuous-time inventory models. *Naval Research Logistics Quarterly*, 33(4):763–774, 1986.