



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Piecewise linear approximations for the static-dynamic uncertainty strategy in stochastic lot-sizing

Citation for published version:

Rossi, R, Kilic, OA & Tarim, SA 2015, 'Piecewise linear approximations for the static-dynamic uncertainty strategy in stochastic lot-sizing', *Omega*, vol. 50, no. 1, pp. 126-140.
<https://doi.org/10.1016/j.omega.2014.08.003>

Digital Object Identifier (DOI):

[10.1016/j.omega.2014.08.003](https://doi.org/10.1016/j.omega.2014.08.003)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

Published In:

Omega

Publisher Rights Statement:

© Rossi, R., Kilic, O. A., & Tarim, S. A. (2014). A unified modeling approach for the static-dynamic uncertainty strategy in stochastic lot-sizing. Omega.

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



A unified modeling approach for the static-dynamic uncertainty strategy in stochastic lot-sizing

Roberto Rossi,^{*,1} Onur A. Kilic,² S. Armagan Tarim,²

¹Business School, University of Edinburgh, Edinburgh, UK
roberto.rossi@ed.ac.uk

²Institute of Population Studies, Hacettepe University, Ankara, Turkey
{onuralp,armagan.tarim}@hacettepe.edu.tr

Abstract

In this paper, we develop mixed integer linear programming models to compute near-optimal policy parameters for the non-stationary stochastic lot sizing problem under Bookbinder and Tan's static-dynamic uncertainty strategy. Our models build on piecewise linear upper and lower bounds of the first order loss function. We discuss different formulations of the stochastic lot sizing problem, in which the quality of service is captured by means of backorder penalty costs, non-stockout probability, or fill rate constraints. These models can be easily adapted to operate in settings in which unmet demand is backordered or lost. The proposed approach has a number of advantages with respect to existing methods in the literature: it enables seamless modelling of different variants of the above problem, which have been previously tackled via ad-hoc solution methods; and it produces an accurate estimation of the expected total cost, expressed in terms of upper and lower bounds. Our computational study demonstrates the effectiveness and flexibility of our models.

keywords: stochastic lot sizing; static-dynamic uncertainty; first order loss function; non-stockout probability; fill rate; penalty cost; piecewise linearisation

Corresponding author: Roberto Rossi, University of Edinburgh Business School, EH8 9JS, Edinburgh, United Kingdom, phone: +44(0)131 6515239, email: roberto.rossi@ed.ac.uk

1 Introduction

Consider the non-stationary stochastic lot sizing problem — the stochastic extension of the well-known dynamic lot sizing problem [43]. This is a finite-horizon periodic review single-item single-stocking location inventory control problem in which demand is stochastic and non-stationary. Bookbinder and Tan [6] discuss three main control strategies that can be adopted in stochastic lot sizing problem: static, static-dynamic, and dynamic uncertainty. The static uncertainty strategy is rather conservative. The decision maker determines both timing and size of orders at the very beginning of the planning horizon. A less conservative strategy is the static-dynamic uncertainty, in which inventory reviews are fixed at the beginning of the planning horizon, while associated order quantities are decided upon only when orders are issued. The dynamic uncertainty strategy allows the decision maker to decide dynamically at each time period whether or not to place an order and how much to order. Each of these strategies has different advantages and disadvantages. For instance, the dynamic uncertainty strategy is known to be cost-optimal [25]. The static uncertainty is appealing in material requirement planning systems, for which order synchronisation is a key concern [38]. The static-dynamic uncertainty strategy has advantages in organising joint replenishments and shipment consolidation [15, 29].

In this study, we focus our attention on the static-dynamic uncertainty strategy, which offers a stable replenishment plan while effectively hedging against uncertainty [13, 42]. An important question regarding the static-dynamic uncertainty strategy is how to determine order quantities at inventory review periods when a replenishment schedule is given. In this context, Özen et al. [16] showed that it is optimal to determine order quantities by means of an order-up-to policy. This result leads to the following characterisation of the static-dynamic uncertainty strategy: at each review period, the decision maker observes the actual inventory position (i.e. on-hand inventory, plus outstanding orders, minus backorder) and places an order so as to increase the inventory position up to a given order-up-to level. Key decisions for the static-dynamic uncertainty strategy include an inventory review schedule and an order-up-to level for each review period — these decisions must be fixed at the beginning of the planning horizon.

We build on recently introduced piecewise linear upper and lower bounds for the first order loss function and its complementary function [24], which are based on distribution independent bounding techniques from stochastic programming: Jensen’s and Edmundson-Madanski’s inequalities [12, p. 167-168]. In contrast to earlier works in the literature, we show that these bounds can be used to estimate inventory holding costs, backorder costs and/or service levels, and that they translate into readily available lower and upper bounds on the optimal expected total costs. Furthermore, for the special case in which demand is normally distributed,

the model relies on standard linearisation parameters provided in [24].

Our analysis leads to a unified modelling approach that captures several variants of the problem and that is based on standard mixed-integer linear programming models. Some of these variants have been previously addressed in the literature, whereas some other have not. More specifically, we consider different assumptions on the way unsatisfied demand is modelled: backorder and lost sales. We also consider different service quality measures commonly employed in the inventory control literature [see e.g. 29, pp. 244—246]: penalty cost per unit short per period, non-stockout probability¹ (α service level), cycle fill rate² (β^{cyc} service level), and fill rate³ (β service level).

Our contributions to the inventory control literature are the following:

- we develop enhanced MILP formulations that enable seamless modelling of the non-stationary stochastic lot sizing problem under each of the four measures of service quality discussed;
- in contrast to other approaches in the literature our MILP models bound from above and below the cost of an optimal plan by using a piecewise linear approximation of the loss function; by increasing the number of segments, precision can be improved ad libitum;
- we discuss how to build these MILP models for the case in which demand in each period follows a generic probability distribution;
- for the special case in which demand in each period is normally distributed, we demonstrate how the MILP formulations can be conveniently constructed via standard linearisation coefficients;
- we discuss for the first time in the literature how to handle the case in which demand that occurs when the system is out of stock is lost, i.e. lost sales;
- we discuss the first MILP formulation in the literature that captures the case in which service quality is modelled using a standard β service level in line with the definition found in many textbooks on inventory control, such as Axsäter [3], Hadley and Whitin [9], Silver et al. [29].
- we present an extensive computational study to show that (i) the optimality gap shrinks exponentially fast as the number of segments in the piecewise

¹a lower bound on the non-stockout probability in any period over the planning horizon

²a lower bound on the expected fraction of demand that is routinely satisfied from stock for each replenishment cycle

³a lower bound on the expected fraction of demand that is routinely satisfied from stock over the planning horizon

linearisation increases, (ii) the number of segments adopted only marginally affects computational efficiency.

2 Literature survey

Due to its practical relevance, a large body of literature has emerged on the static-dynamic uncertainty strategy over the last decades. Here, we review some key contributions which are of particular importance in the context of the current work, and reflect upon our contribution.

To keep our discussion focused, since all the models we will discuss in the following sections operate under a static-dynamic uncertainty strategy, in our literature review we do not survey works related to the static uncertainty [see e.g. 10, 37, 39, 40] and to the dynamic uncertainty strategy, [see e.g. 5].

Early works on the stochastic lot sizing problem concentrated on easy-to-compute heuristics. Silver [26] and Askin [1] studied the problem under penalty costs, and proposed simple heuristics based on the least period cost method. These heuristics can be regarded as stochastic extensions of the well-known Silver-Meal heuristic [28].

Bookbinder and Tan [6] studied the problem under α service level constraints and introduced the terminology “static uncertainty,” “dynamic uncertainty,” and “static-dynamic uncertainty.” They developed a method that sequentially determines the timing of replenishments and corresponding order-up-to levels for the static-dynamic uncertainty strategy. Following this seminal work, a variety of further studies — which significantly differ in terms of underlying service quality measures and modeling approaches — aimed to determine the optimal replenishment schedule and order-up-to levels simultaneously under Bookbinder and Tan’s static-dynamic uncertainty strategy.

Tarim and Kingsman [31] discussed the first MILP formulation under α service level constraints. In contrast to [6], this formulation simultaneously determines the replenishment schedule and corresponding order-up-to levels. Efficient reformulations operating under the same assumptions were discussed in [21, 35]. [21] developed a state space augmentation approach; while [35] implemented a branch and bound algorithm. In addition, Constraint Programming reformulations based on a novel modelling tool, i.e. global chance constraints, were discussed in [33, 34]. Finally, an exact, although computationally intensive, Constraint Programming approach was discussed in [18]. Extensions to the case of a stochastic delivery lead time were discussed in [19, 20].

Tarim and Kingsman [32] developed the first MILP formulation for the case in which service quality is modelled using a penalty cost scheme. Rossi et al. [23] discussed an efficient Constraint Programming reformulation exploiting optimization

oriented global stochastic constraints.

Özen et al. [16] discussed a dynamic programming solution algorithm and two ad-hoc heuristics named “approximation” and “relaxation” heuristics, respectively; the authors analyse both the penalty cost and the α service level cases. The “approximation” heuristic operates under the assumption that scenarios in which the actual stock exceeds the order-up-to-level for a given review are negligible and can be safely ignored; while the “relaxation” heuristics operates by relaxing those constraints in Tarim and Kingsman’s model that force order sizes in each period to be nonnegative.

Tempelmeier [36] introduced an MILP formulation for the case in which service quality is modelled via β^{cyc} service level constraints.

A key issue in all the aforementioned studies is the computation of the true values of expected on-hand inventories and stock-outs, and thereby associated costs and/or service levels. These values can only be derived from the (complementary) first-order loss function of the demand [see 30, p. 338] — a non-linear function that cannot be readily embedded into the proposed MILP models.

Tarim and Kingsman [31] — but also [18, 19, 20, 21, 33, 34, 35] — bypassed the issue by approximating the expected on-hand inventory by the expected inventory position. This approach could work well for inventory systems that operate under a very high non-stockout probability. However, it may result in highly sub-optimal solutions when the probability of observing a stock-out is not negligible.

Tarim and Kingsman [32] used a piecewise linear approximation of the standard loss function to approximate expected holding and penalty costs. The piecewise linear function is fitted to the nonlinear cost function by using an approach that minimises the maximum absolute approximation error. The power of this approach is that the piecewise linearisation is based on standard linearisation coefficients that can be computed once and then reused for any normally distributed demand. Unfortunately, the approximation proposed may either over or underestimate the original cost therefore it becomes hard to assess how far a given solution may be from the true optimal one, i.e. its optimality gap. Furthermore, this approximation is not easily extended to demands following a generic distribution or to models operating under service level measures rather than a penalty cost scheme.

Özen et al. [16] approximations require ad-hoc algorithms and cannot be easily extended to handle β service level constraints or lost sales.

Tempelmeier [36] tabulates the complementary first order loss function and then uses binary variables to retrieve the holding cost associated with a given replenishment plan. A similar tabulation is employed to enforce the prescribed β^{cyc} service level. However, this tabulation is carried out by considering each possible replenishment cycle⁴ independently. This strategy disregards cost and service level

⁴A replenishment cycle is the time interval between two successive inventory reviews.

dependencies that may exist among successive replenishment cycles. For this very same reason, it cannot be employed to model classic β service level constraints. To the best of our knowledge, no formulation exists in the literature for the case in which service quality is modelled via standard β service level constraints [3, 9, 29].

The issue of computing expected on-hand inventories and stock-outs is topical in inventory control, as witnessed by a number of recent works [see e.g. 2, 27]. All approaches surveyed above address particular instances of the problem. The contribution of this paper is unique and novel in the sense that it introduces a unified modelling approach for static-dynamic uncertainty strategy based on linear approximations of the first order loss function. Furthermore, this unified modelling approach can be used to address the issue of computing static-dynamic uncertainty policy parameters under lost sales, which has not been addressed yet in the literature.

3 Piecewise linearisation of loss functions

For convenience, a list of all symbols used in the rest of the paper is provided in Appendix I. Consider a random variable ω with expected value $\tilde{\omega}$ and a scalar variable x . The first order loss function is defined as $\mathcal{L}(x, \omega) = \text{E}[\max(\omega - x, 0)]$, where E denotes the expected value. The complementary first order loss function is defined as $\widehat{\mathcal{L}}(x, \omega) = \text{E}[\max(x - \omega, 0)]$. It is known that there is a close relationship between these two functions, as stated in the following lemma.

Lemma 1 ([30] p. 338, C.5).

$$\mathcal{L}(x, \omega) = \widehat{\mathcal{L}}(x, \omega) - (x - \tilde{\omega}). \quad (1)$$

The first order loss function and its complementary function play a key role in inventory models, since they are essential to compute expected holding and penalty costs, as well as a number of service measure such as the β^{cyc} and β service levels.

A common approach in computing near-optimal control parameters of the static-dynamic uncertainty strategy is to formulate the problem as a certainty equivalent MILP. Unfortunately, the loss function is non-linear and cannot be easily embedded in MILP models. To overcome this issue, we adopt a piecewise linearisation approach similar to the one recently discussed in Rossi et al. [24]. This approach is based on classical inequalities from stochastic programming: Jensen's and Edmundson-Madanski inequalities [12, p. 167-168] and can be applied to random variables following a generic probability distribution.

It is known that both the first order loss function $\mathcal{L}(x, \omega)$ and its complementary function $\widehat{\mathcal{L}}(x, \omega)$ are convex in x regardless of the distribution of ω . For this reason, both Jensen's lower bound and Edmundson-Madanski upper bound are applicable

to the first order loss function and its complementary function. More formally, let $g_\omega(\cdot)$ denote the probability density function of ω and consider a partition of the support Ω of ω into W disjoint compact subregions $\Omega_1, \dots, \Omega_W$. We define, for all $i = 1, \dots, W$

$$p_i = \Pr\{\omega \in \Omega_i\} = \int_{\Omega_i} g_\omega(t) dt \quad \text{and} \quad E[\omega|\Omega_i] = \frac{1}{p_i} \int_{\Omega_i} t g_\omega(t) dt \quad (2)$$

Lemma 2. *For the complementary first order loss function the lower bound $\widehat{\mathcal{L}}_{lb}(x, \omega)$, where*

$$\widehat{\mathcal{L}}(x, \omega) \geq \widehat{\mathcal{L}}_{lb}(x, \omega) = \sum_{i=1}^W p_i \max(x - E[\omega|\Omega_i], 0)$$

is a piecewise linear function with $W + 1$ segments. The i -th linear segment of $\widehat{\mathcal{L}}_{lb}(x, \omega)$ is

$$\widehat{\mathcal{L}}_{lb}^i(x, \omega) = x \sum_{k=1}^i p_k - \sum_{k=1}^i p_k E[\omega|\Omega_k] \quad E[\omega|\Omega_i] \leq x \leq E[\omega|\Omega_{i+1}], \quad (3)$$

where $i = 1, \dots, N$; furthermore, the 0-th segment is $x = 0$, $-\infty \leq x \leq E[\omega|\Omega_1]$.

This lower bound is a direct application of Jensen's inequality. Let then e_W denote the maximum approximation error for the lower bound in Lemma 2 associated with a given partition comprising W regions. A piecewise linear upper bound, i.e. Edmundson-Madanski's bound, can be obtained by shifting up the lower bound in Lemma 2 by a value e_W .

Lemma 3. *For the complementary first order loss function the upper bound $\widehat{\mathcal{L}}_{ub}(x, \omega)$, where*

$$\widehat{\mathcal{L}}(x, \omega) \leq \widehat{\mathcal{L}}_{ub}(x, \omega) = \sum_{i=1}^W p_i \max(x - E[\omega|\Omega_i], 0) + e_W$$

is a piecewise linear function with $W + 1$ segments. The i -th linear segment of $\widehat{\mathcal{L}}_{ub}(x, \omega)$ is

$$\widehat{\mathcal{L}}_{ub}^i(x, \omega) = x \sum_{k=1}^i p_k - \sum_{k=1}^i p_k E[\omega|\Omega_k] + e_W \quad E[\omega|\Omega_i] \leq x \leq E[\omega|\Omega_{i+1}],$$

where $i = 1, \dots, N$; furthermore, the 0-th segment is $x = e_W$, $-\infty \leq x \leq E[\omega|\Omega_1]$.

Having established these two results, we must then decide how to partition the support Ω in order to obtain good bounds. A number of works discussed how to

obtain an optimal partitioning of the support under a framework that minimises the maximum approximation error [8, 11]. In short, these works demonstrate that, in order to minimise the maximum approximation error, one must find parameters ensuring approximation errors at piecewise function breakpoints are all equal. This result unfortunately does not hold when optimal linearisation parameters must be found for a set of random variables.

Consider a set of random variables $\omega_1, \dots, \omega_n, \dots, \omega_N$ and associated complementary first order loss functions $\widehat{\mathcal{L}}(x, \omega_1), \dots, \widehat{\mathcal{L}}(x, \omega_N)$. From (2) it is clear that, once all p_i have been fixed, all $E[\omega_n|\Omega_k]$ are uniquely determined. The particular structure of (3) makes it, in principle, possible to compute standard p_i coefficients for the whole set of random variables and then select the $E[\omega_n|\Omega_k]$ for a specific $\widehat{\mathcal{L}}_{lb}(x, \omega_n)$ via a binary selector variable y_n , that is

$$\widehat{\mathcal{L}}_{lb}^i(x, \omega) = x \sum_{k=1}^i p_k - \sum_{k=1}^i p_k E[\omega_n|\Omega_k] y_n \quad E[\omega_n|\Omega_i] \leq x \leq E[\omega_n|\Omega_{i+1}] \quad 1 \leq n \leq N \quad (4)$$

where $\sum_{n=1}^N y_n = 1$. These expressions generalise those discussed in [24], which only hold for normally distributed random variables, and they are particularly useful, as we will see in Section 4, if one wants to develop an MILP model for computing Bookbinder and Tan's static-dynamic uncertainty policy parameters.

Unfortunately, computing probability masses p_1, \dots, p_W that minimise the maximum approximation error over a set of random variables is a challenging task. As shown in [17] this is not a problem of convex optimisation as the one faced while computing optimal linearisation parameters for a single loss function. In this work, we will adopt two approximate strategies to compute good probability masses p_1, \dots, p_W : a simple strategy that splits the support into W regions with uniform probability mass; and a more refined local search strategy that uses a combination of simple random sampling and coordinate descent. Implementation details of these heuristic approaches are discussed in Appendix II.

4 Enhanced MILP reformulations

In this section we demonstrate how the results presented so far can be used to derive enhanced a mixed integer programming formulations for the stochastic lot sizing problem under static-dynamic uncertainty strategy. First, we introduce the original stochastic programming formulation of the problem (Section 4.1). We then present MILP models for the problem under α service level constraints (Section 4.2), a penalty cost scheme (Section 4.3), β^{cyc} (Section 4.4) and β (Section 4.5) service level constraints. Finally, we discuss how to extend these models to a lost sales setting. For a complete overview of the models presented the reader may

refer to Appendix III.

4.1 Stochastic lot-sizing

The stochastic programming formulation of the non-stationary stochastic lot-sizing problem was originally presented in [6, pp. 1097–1098]. The formal problem definition is as follows. Customer demand d_t in each period $t = 1, \dots, N$ is a random variable with probability density function $g_t(\cdot)$ and cumulative distribution function $G_t(\cdot)$. There are fixed and variable replenishment costs: the fixed cost is a per order; the variable cost is v per unit ordered. Negative orders are not allowed. A holding cost of h is paid of each unit of inventory carried from one period to the next. I_0 denotes the initial inventory level. Delivery lead-time is not incorporated in the model. When a stockout occurs, all demand is backordered and filled as soon as an adequate supply arrives. There is a service level constraint enforcing a non-stockout probability of at least α in each period — this is known in the inventory control literature as “ α service level” constraint [29]. They finally also assume that the service level is set to a high value, i.e. $\alpha > 0.9$ in order to incorporate management’s perception of the cost of backorders, so that shortage costs can be safely ignored. The objective is to minimise the expected total cost $E[\text{TC}]$, which comprises fixed/variable ordering and holding costs. The resulting model is presented in Fig. 1. In this model I_t represents the inventory level at the end of a period; δ_t takes value 1 if an order is placed in period t ; and Q_t represents the order quantity in period t . Constraints (6) are the inventory conservation constraints: inventory level at the end of period t must be equal to the initial inventory I_0 , plus all order received, minus all demand realised up to period t , since we assume — in line with the original model in the literature — that inventory cannot be disposed or returned to the supplier; constraints (7) set δ_t to one if an order is placed in period t ; finally, (8) enforce the prescribed service level in each period.

The above model can be easily modified to accommodate a penalty cost scheme, in place of the original α service level constraints, [see e.g. 32]. All one has to do is to drop (8) and to replace the original objective function with the following one

$$E[\text{TC}] = \min \int_{d_1} \int_{d_2} \dots \int_{d_N} \sum_{t=1}^N (a\delta_t + h \max(I_t, 0) + b \max(-I_t, 0) + vQ_t) \times \quad (10)$$

$$g_1(d_1)g_2(d_2) \dots g_N(d_N) d(d_1)d(d_2) \dots d(d_N)$$

where b denotes the penalty cost per unit short per period. β^{cyc} and β service level formulations are obtained by replacing (8) with an alternative service measure.

$$\begin{aligned} E[\text{TC}] = \min \int_{d_1} \int_{d_2} \dots \int_{d_N} \sum_{t=1}^N (a\delta_t + h \max(I_t, 0) + vQ_t) \times \\ g_1(d_1)g_2(d_2) \dots g_N(d_N) d(d_1)d(d_2) \dots d(d_N) \end{aligned} \quad (5)$$

subject to, for $t = 1, \dots, N$

$$I_t = I_0 + \sum_{i=1}^t (Q_i - d_i) \quad (6)$$

$$\delta_t = \begin{cases} 1 & \text{if } Q_t > 0, \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

$$\Pr\{I_t \geq 0\} \geq \alpha \quad (8)$$

$$Q_i \geq 0, \delta_t \in \{0, 1\} \quad (9)$$

Figure 1: Stochastic programming formulation of the non-stationary stochastic lot-sizing problem.

4.2 α service level constraints

We now consider the mixed integer programming formulation of Tarim and Kingsman [31] for computing near-optimal inventory control policy parameters under Bookbinder and Tan's static-dynamic uncertainty strategy. According to this strategy, inventory review times as well as their respective order-up-to-levels must be all fixed at the beginning of the planning horizon. However, actual order quantities are determined only after demand has been observed.

In what follows, M denotes a very large number and \tilde{x} denotes the expected value of x . Tarim and Kingsman's model is presented in Fig. 2. This certainty equivalent model comprises three sets of decision variables: \tilde{I}_t , representing the expected closing inventory level at the end of period t ; δ_t , a binary variable representing the inventory review decision at period t ; and P_{jt} , a binary variable which is set to one if and only if the most recent inventory review before period t was carried out in period j . By observing that, for a period t in which an order is placed (i.e. $\delta_t = 1$) the order-up-to-level S_t is simply $S_t = \tilde{I}_t + \tilde{d}_t$, it follows that by solving the above model policy parameters are immediately obtained.

Constraints in the certainty equivalent model neatly reflect those in the original stochastic programming model. More specifically, (12) enforces the inventory conservation constraints; (13) is the reordering condition; and (14), (15), and (16)

$$\mathbb{E}[\text{TC}] = -vI_0 + v \sum_{t=1}^N \tilde{d}_t + \min \sum_{t=1}^N (a\delta_t + h\tilde{I}_t) + v\tilde{I}_N \quad (11)$$

subject to, for $t = 1, \dots, N$

$$\tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1} \geq 0 \quad (12)$$

$$\tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1} \leq \delta_t M \quad (13)$$

$$\tilde{I}_t \geq \sum_{j=1}^t \left(G_{d_{jt}}^{-1}(\alpha) - \sum_{k=j}^t \tilde{d}_k \right) P_{jt} \quad (14)$$

$$\sum_{j=1}^t P_{jt} = 1 \quad (15)$$

$$P_{jt} \geq \delta_j - \sum_{k=j+1}^t \delta_k \quad j = 1, \dots, t \quad (16)$$

$$P_{jt} \in \{0, 1\} \quad j = 1, \dots, t \quad (17)$$

$$\delta_t \in \{0, 1\} \quad (18)$$

Figure 2: MILP formulation of the non-stationary stochastic lot-sizing problem under the static-dynamic uncertainty strategy [31].

enforce the prescribed service level α . In (14), $G_{d_{jt}}^{-1}(\alpha)$ denotes the α -quantile of the inverse cumulative distribution function of the random variable $d_j + \dots + d_t$. Finally, the objective function is obtained by observing

$$\mathbb{E}\left[v \sum_{t=1}^N Q_t\right] = -vI_0 + v \sum_{t=1}^N \tilde{d}_t + v\tilde{I}_N, \quad (19)$$

where $-vI_0 + v \sum_{t=1}^N \tilde{d}_t$ is a constant.

Following an assumption originally introduced by Bookbinder and Tan [6], Tarim and Kingsman [31] approximate the holding cost component in the original objective function, which we recall was $\mathbb{E}[\max(I_t, 0)]$, via the expression $h\tilde{I}_t$. To overcome this limitation of the model, we introduce two new sets of decision variables: $\tilde{I}_t^{lb} \geq 0$ and $\tilde{I}_t^{ub} \geq 0$ for $t = 1, \dots, N$, which represent, respectively, a lower and an upper bound to the true value of $\mathbb{E}[\max(I_t, 0)]$. The objective function

then can be rewritten as

$$\mathbb{E}[\text{TC}] = -vI_0 + v \sum_{t=1}^N \tilde{d}_t + \min \sum_{t=1}^N (a\delta_t + h\tilde{I}_t^{lb}) + v\tilde{I}_N \quad (20)$$

if our aim is to compute a lower bound for the cost of an optimal plan, or as

$$\mathbb{E}[\text{TC}] = -vI_0 + v \sum_{t=1}^N \tilde{d}_t + \min \sum_{t=1}^N (a\delta_t + h\tilde{I}_t^{ub}) + v\tilde{I}_N \quad (21)$$

if our aim is to compute an upper bound for the cost of an optimal plan. We next discuss how to constrain \tilde{I}_t^{lb} and \tilde{I}_t^{ub} . Our discussion applies to demand d_k , for $k = 1, \dots, N$ following a generic distribution. Consider the random variable d_{jt} representing the convolution $d_j + \dots + d_t$. We select a priori a number W of adjacent regions Ω_i into which the support of d_{jt} must be partitioned. As discussed, this partitioning will produce a piecewise linear approximation comprising $W + 1$ segments. We also fix a priori the probability mass $p_i = \Pr\{d_{jt} \in \Omega_i\}$ that must be associated with each region Ω_i . As discussed in Section 3, there are several possible strategies to assign probability masses p_i to regions. For instance, we may ensure uniform coverage, i.e. all region must have the same probability mass, or we may select regions — by using a heuristic or exact approach — in order to minimise the maximum approximation error over all possible convolutions d_{jt} , for $t = 1, \dots, N$ and $j = 1, \dots, t$. Regardless of the strategy we adopt, once all p_i are known, regions Ω_i are uniquely determined and the associated conditional expectation $\mathbb{E}[d_{jt}|\Omega_i]$ can be immediately computed using off-the-shelf software. Finally, the maximum approximation error e_W^{jt} associated with the linearisation of $\hat{\mathcal{L}}(x, d_{jt})$ can be found by checking the linearisation error at the W possible breakpoints of the piecewise linear function obtained. Having precomputed all these values for $t = 1, \dots, N$ and $j = 1, \dots, t$, we introduce the following constraints in the model

$$\tilde{I}_t^{lb} \geq (\tilde{I}_t + \sum_{j=1}^t \tilde{d}_{jt}P_{jt}) \sum_{k=1}^i p_k - \sum_{j=1}^t \left(\sum_{k=1}^i p_k \mathbb{E}[d_{jt}|\Omega_i] \right) P_{jt} \quad t = 1, \dots, N; \quad i = 1, \dots, W \quad (22)$$

This expression follows from Lemma 2 and closely resembles (4). Consider a replenishment in period j covering periods j, \dots, t with associated order-up-to-level S . Our aim is to enforce $\tilde{I}_t^{lb} \geq \hat{\mathcal{L}}_{lb}^i(S, d_{jt})$ for $i = 1, \dots, W$, since \tilde{I}_t^{lb} represents a lower bound for the expected positive inventory at the end of period t . By observing that $S = \tilde{I}_t + \tilde{d}_{jt}$, we obtain the above expression. We then obtain \tilde{I}_t^{ub} from \tilde{I}_t , by noting that a piecewise linear upper bound can be derived by adding

the maximum estimation error to the Jensen's piecewise linear lower bound [24].

$$\tilde{I}_t^{ub} \geq \left(\tilde{I}_t + \sum_{j=1}^t \tilde{d}_{jt} P_{jt} \right) \sum_{k=1}^i p_k + \sum_{j=1}^t \left(e_W^{jt} - \sum_{k=1}^i p_k \mathbb{E}[d_{jt} | \Omega_i] \right) P_{jt} \quad \begin{array}{l} t = 1, \dots, N, \\ i = 1, \dots, W; \end{array} \quad (23)$$

where $\tilde{I}_t^{ub} \geq \sum_{j=1}^t e_W^{jt} P_{jt}$ for $t = 1, \dots, N$. The special case in which demand in each period follows a normal distribution is discussed in Appendix IV.

4.3 Penalty cost scheme

The model discussed in Section 4.2 can be easily modified to accommodate a penalty cost b per unit short per period in place of the α service level constraints discussed in Tarim and Kingsman [31]. This revised model resembles the one discussed in Tarim and Kingsman [32]. However, as discussed in the previous section, our formulation is more accurate, because the expected total cost of a plan can be now bounded from above and below. It is also more general, since the discussion in Tarim and Kingsman [32] is limited to normally distributed demand.

In the new model, we introduce two new sets of variables $\tilde{B}_t^{lb} \geq 0$ and $\tilde{B}_t^{ub} \geq 0$ for $t = 1, \dots, N$, which represent a lower and upper bound, respectively, for the true value of $\mathbb{E}[-\min(I_t, 0)]$ and thus allow us to compute lower and upper bounds for the expected backorders in each period. The objective function then becomes

$$\mathbb{E}[\text{TC}] = -vI_0 + v \sum_{t=1}^N \tilde{d}_t + \min \sum_{t=1}^N (a\delta_t + h\tilde{I}_t^{lb} + b\tilde{B}_t^{lb}) + v\tilde{I}_N \quad (24)$$

if our aim is to compute a lower bound for the cost of an optimal plan, or

$$\mathbb{E}[\text{TC}] = -vI_0 + v \sum_{t=1}^N \tilde{d}_t + \min \sum_{t=1}^N (a\delta_t + h\tilde{I}_t^{ub} + b\tilde{B}_t^{ub}) + v\tilde{I}_N \quad (25)$$

if our aim is to compute an upper bound for the cost of an optimal plan. Finally, we must remove constraints (14), since we are operating under a penalty cost scheme and not under a service level constraints.

Once more, we assume that demand in each period follows a generic distribution; we obtain \tilde{B}_t^{lb} and \tilde{B}_t^{ub} from \tilde{I}_t by exploiting Lemma 1.

$$\tilde{B}_t^{lb} \geq -\tilde{I}_t + \left(\tilde{I}_t + \sum_{j=1}^t \tilde{d}_{jt} P_{jt} \right) \sum_{k=1}^i p_k - \sum_{j=1}^t \left(\sum_{k=1}^i p_k \mathbb{E}[d_{jt} | \Omega_i] \right) P_{jt} \quad \begin{array}{l} t = 1, \dots, N, \\ i = 1, \dots, W; \end{array} \quad (26)$$

where $\tilde{B}_t^{ub} \geq -\tilde{I}_t$ and

$$\tilde{B}_t^{ub} \geq -\tilde{I}_t + \left(\tilde{I}_t + \sum_{j=1}^t \tilde{d}_{jt} P_{jt} \right) \sum_{k=1}^i p_k + \sum_{j=1}^t \left(e_W^{jt} - \sum_{k=1}^i p_k \mathbb{E}[d_{jt} | \Omega_i] \right) P_{jt} \quad \begin{array}{l} t = 1, \dots, N, \\ i = 1, \dots, W; \end{array} \quad (27)$$

where $\tilde{B}_t^{ub} \geq -\tilde{I}_t + \sum_{j=1}^t e_W^{jt} P_{jt}$. The case in which demand in each period follows a normal distribution is discussed in Appendix IV.

4.4 β^{cyc} service level constraints

The model discussed in Section 4.3 can be modified to accommodate β^{cyc} service level constraints, defined in [36] as

$$1 - \max_{i=1, \dots, m} \left[\mathbb{E} \left\{ \frac{\text{Total backorders in replenishment cycle } i}{\text{Total demand in replenishment cycle } i} \right\} \right]. \quad (28)$$

β^{cyc} represents a lower bound on the expected fraction of demand that is routinely satisfied from stock for each replenishment cycle. The revised model resembles the one discussed in Tempelmeier [36]. However, we aim to stress that in this latter work the author enforces the prescribed service level by precomputing order-up-to-levels and cycle holding costs in a table, rather than using a piecewise linearisation of the loss function as we do. One of the advantages of our approach is that it is able to account for dependencies among opening stock levels and service levels of consecutive cycles. As we will show in Section 4.5, the ability to capture these dependencies is important if we aim to generalise the model to a more classical definition of “fill rate”.

The discussion below is distribution independent; we modify the model discussed in Section 4.3 by introducing service level constraints

$$\tilde{B}_t^{lb} \leq (1 - \beta^{\text{cyc}}) \sum_{j=1}^t P_{jt} \tilde{d}_{jt} \quad t = 1, \dots, N, \quad (29)$$

if our aim is to compute a lower bound for the cost of an optimal plan; or

$$\tilde{B}_t^{ub} \leq (1 - \beta^{\text{cyc}}) \sum_{j=1}^t P_{jt} \tilde{d}_{jt} \quad t = 1, \dots, N, \quad (30)$$

if our aim is to compute an upper bound for the cost of an optimal plan. Constraints (29) and (30) directly follow from (28). Finally, the objective function is (20) if our aim is to compute a lower bound for the cost of an optimal plan, or

(21) if our aim is to compute an upper bound.

4.5 β service level constraints

The model discussed in Tempelmeier [36] captures a definition of fill rate that is not conventional in the inventory literature. This issue has been discussed in Rossi et al. [22]. To date, no modelling strategy exists for the conventional fill rate under a static-dynamic uncertainty control policy. In this section, we introduce an alternative MILP reformulation that captures a definition of β service level that is in line with the definition found in many textbooks on inventory control Axsäter [3], Hadley and Whitin [9], Silver et al. [29].

In Axsäter [3], the author defines β service level as “the expected fraction of demand satisfied immediately from stock on hand”. In the context of finite horizon inventory models, e.g. Chen et al. [7], Thomas [41] this definition is formalized as

$$1 - \mathbb{E} \left\{ \frac{\text{Total backorders within the planning horizon}}{\text{Total demand within the planning horizon}} \right\}, \quad (31)$$

The static-dynamic uncertainty strategy divides the finite planning horizon into a number, say m , of consecutive replenishment cycles. We can re-write (31) by taking these into account as

$$1 - \mathbb{E} \left\{ \frac{\sum_{i=1}^m \text{Total backorders within the } i\text{'th replenishment cycle}}{\sum_{i=1}^m \text{Total demand within the } i\text{'th replenishment cycle}} \right\}. \quad (32)$$

However, in [36] the same β service level is imposed on each and every cycle within the planning horizon. Clearly (32) is different from (28): the original definition imposes a β service level throughout the planning horizon, whereas the definition in [36] imposes a β service level on each replenishment cycle within the planning horizon independently. This latter definition is thus more restrictive. Adopting the original definition may therefore bring substantial cost benefits. To illustrate this, one may consider instances discussed in [36, p. 191], for which the cost reduction with respect to policies obtained via the model discussed in Section 4.4 ranges from 0.4% to 6.4%.

We modify as follows the model in Section 4.4 to implement the classical measure outlined in (32). Also in this cases the changes discussed below are distribution independent. We introduce two new set of nonnegative variables \tilde{C}_t^{lb} and \tilde{C}_t^{ub} for $t = 0, \dots, N$. These variables express a lower and an upper bound, respectively, to the expected total backorders within the replenishment cycle that ends at period t , if any exists. Hence, \tilde{C}_t^{lb} (resp. \tilde{C}_t^{ub}) should be equal to \tilde{B}_t^{lb} (resp. \tilde{B}_t^{ub}), if t is the last period of a replenishment cycle; otherwise \tilde{C}_t^{lb} (resp. \tilde{C}_t^{ub}) should be equal

to 0. We enforce this fact as follows:

$$\tilde{C}_t^{lb} \geq \tilde{B}_t^{lb} - (1 - \delta_{t+1}) \sum_{k=1}^t \tilde{d}_k \quad t = 0, \dots, N - 1, \quad (33)$$

$$\tilde{C}_t^{ub} \geq \tilde{B}_t^{ub} - (1 - \delta_{t+1}) \sum_{k=1}^t \tilde{d}_k \quad t = 0, \dots, N - 1. \quad (34)$$

where $\tilde{B}_0^{lb} = \tilde{B}_0^{ub} = \tilde{C}_0^{lb} = \tilde{C}_0^{ub} = I_0$. Finally, we must ensure that $\tilde{C}_N^{lb} = \tilde{B}_N^{lb}$ and $\tilde{C}_N^{ub} = \tilde{B}_N^{ub}$. We then use these new variables to build constraint

$$\sum_{t=1}^N \tilde{C}_t^{lb} \leq (1 - \beta) \sum_{t=1}^N \tilde{d}_t \quad (35)$$

which will replace (29), if our aim is to compute a lower bound for the cost of an optimal plan; and constraint

$$\sum_{t=1}^N \tilde{C}_t^{ub} \leq (1 - \beta) \sum_{t=1}^N \tilde{d}_t \quad (36)$$

which will replace (30), if our aim is to compute an upper bound. Constraints 35 and 36 directly follow from 32.

5 Lost sales

In this section we briefly sketch the extension of the models discussed in the previous section to the case in which demand that occurs when the system is out of stock is lost, i.e. lost sales setting. The discussion is purportedly short, since the derived models are quite similar to those already presented and there are only few adjustments that are necessary to adapt our models to this new settings.

Under lost sales we need to take into account the fact that if inventory drops to zero, demand that occurs until the next order arrives will not be met. Under a static-dynamic uncertainty control policy, this means the actual order quantity will never exceed the order-up-to-level.

In this settings it is crucial to set up the model in such a way as to account for the opportunity cost associated with units of demand that are not met by a given control policy. For this reason, we must introduce a new parameter s that represents the selling price of a product; we then let $m = s - v$ be the margin — i.e. unit selling price minus unit ordering cost — for an item sold. The resulting stochastic programming model under lost sales is shown in Fig. 3. In

$$\begin{aligned}
\text{E[TP]} = sI_0 + \max \int_{d_1} \int_{d_2} \dots \int_{d_N} & \left[\sum_{t=1}^N (mQ_t - a\delta_t - h \max(I_t, 0)) - s \max(I_N, 0) \right] \times \\
& g_1(d_1)g_2(d_2) \dots g_N(d_N) d(d_1)d(d_2) \dots d(d_N)
\end{aligned} \tag{37}$$

subject to, for $t = 1, \dots, N$

$$I_t + d_t - I_{t-1} \geq 0 \tag{38}$$

$$Q_t = I_t + d_t - \max(I_{t-1}, 0) \tag{39}$$

$$\delta_t = \begin{cases} 1 & \text{if } Q_t > 0, \\ 0 & \text{otherwise} \end{cases} \tag{40}$$

$$\Pr\{I_t \geq 0\} \geq \alpha \tag{41}$$

$$Q_t \geq 0, \delta_t \in \{0, 1\} \tag{42}$$

Figure 3: Stochastic programming formulation of the non-stationary stochastic lot-sizing problem under lost sales.

this model E[TP] represents the expected total profit, which we aim to maximise. δ_t is a binary decision variable that is set to one if we order items at period t , i.e. constraints (40). Q_t represents the order quantity in period t , which must be greater or equal to zero. I_t is a random variable that represents the inventory level at the end of period t ; despite lost sales, for convenience, we assume that I_t may take negative values. In the objective function, we multiply the margin m by the number of items Q_t ordered in period t ; we then subtract ordering cost a , if an order is placed in period t (i.e. $\delta_t = 1$), and the holding cost h on items that remain in stock at the end of period t . There is a further term $-s \max(I_N, 0)$ to reflect the fact that items in stock at the end of the planning horizon will not be sold and thus the associated selling price s should not be included in the total profit. Term $\max(I_{t-1}, 0)$ in constraints (39) makes sure that the order quantity does not include any lost sale from the previous period. Constraints (38) ensure that the inventory level at the end of period t is greater or equal to the inventory level at the end of period $t - 1$ plus the realised demand in period t ; this makes sure that items in stock in a period and not sold are brought to the next period. Constraints (41) enforce an α service level and can be easily replaced by other service measures such as a β^{cyc} or a β service level. A formulation under a penalty cost scheme is also easily obtained by removing the service level constraints and

by modifying the objective function as illustrated in previous sections.

An MILP formulation of the problem in Fig. 3 under the static-dynamic uncertainty strategy is shown in Fig. 4. In this model, by taking expectations,

$$\mathbb{E}[\text{TP}] = sI_0 + \max \sum_{t=1}^N (m\tilde{Q}_t - a\delta_t - h\tilde{I}_t^{lb}) - s\tilde{I}_N^{lb} \quad (43)$$

subject to, for $t = 1, \dots, N$

$$\tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1} \geq 0 \quad (44)$$

$$\tilde{Q}_t \geq \tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1}^{lb} + (1 - \delta_t)M \quad (45)$$

$$\tilde{Q}_t \leq \tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1}^{lb} - (1 - \delta_t)M \quad (46)$$

$$\tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1} \leq \tilde{Q}_t M \quad (47)$$

$$\tilde{Q}_t \leq \delta_t M \quad (48)$$

$$\tilde{I}_t \geq \sum_{j=1}^t \left(G_{d_{jt}}^{-1}(\alpha) - \sum_{k=j}^t \tilde{d}_k \right) P_{jt} \quad (49)$$

$$\sum_{j=1}^t P_{jt} = 1 \quad (50)$$

$$P_{jt} \geq \delta_j - \sum_{k=j+1}^t \delta_k \quad j = 1, \dots, t \quad (51)$$

constraints (22)

$$P_{jt} \in \{0, 1\} \quad j = 1, \dots, t \quad (52)$$

$$\tilde{Q}_t \geq 0, \delta_t \in \{0, 1\} \quad (53)$$

Figure 4: MILP formulation of the non-stationary stochastic lot-sizing problem under the static-dynamic uncertainty strategy and lost sales.

constraints (38) translate into (44). Constraints (39) translate into (45) and (46). Terms $\max(I_t, 0)$ in the objective function and in constraints (39) can be handled by using an auxiliary variable \tilde{I}_t^{lb} that represents a lower bound for $\mathbb{E}[\max(I_t, 0)]$ computed as before via a piecewise linearisation of the complementary first order loss function, i.e. constraints (22). Constraints (47) ensures that the expected inventory level at the end of period t is greater than the expected inventory level

at the end of period $t - 1$ plus the expected demand in period t if and only if an order has been placed in period t , i.e. $\tilde{Q}_t > 0$. Constraints (40) translate into (48). Finally, service level constraints (41) translate into (49), (50) and (51) following a strategy similar to the one in [31], which we illustrated in Section 4.2. The model presented can be used to compute an upper bound for $E[\text{TP}]$ — note that underestimating buffer stocks, i.e. \tilde{I}_t^{lb} leads to lower holding costs and to an overestimation of the expected order quantity and associated margins $m\tilde{Q}_t$ in the objective function. If we aim to compute a lower bound instead, all occurrences of \tilde{I}_t^{lb} should be replaced by \tilde{I}_t^{ub} and constraints (22) should be replaced by constraints (23). Other MILP formulations under β^{cyc} and β service levels are obtained in a similar fashion, since only the service level constraints of the model are affected by this change. A penalty cost formulation is also easily obtained, by dropping the service level constraints and by modifying the objective function as follows

$$E[\text{TP}] = sI_0 + \max \sum_{t=1}^N (m\tilde{Q}_t - a\delta_t - h\tilde{I}_t^{lb} - b\tilde{B}_t^{lb}) - s\tilde{I}_N^{lb} \quad (54)$$

A detailed overview of these models is given in Appendix III. Finally, a remark that should be made is that the model presented in Section 4.3 charges penalty cost on a “per unit short per unit time” basis. This might be not appropriate under lost sales, since it is common practice to charge penalty cost on a “per unit short” basis under this setting. This problem can be easily overcome by charging the penalty cost b not on \tilde{B}_t^{lb} , which bounds the expected units short at the end of each period, but on \tilde{C}_t^{lb} , which bounds the expected units short at the end of each replenishment cycle — as defined in Section 4.5.

6 Computational experience

In this section we present an extensive computational analysis of the models previously discussed. The experiments below were conducted by using CPLEX 12.3 on a 2.13 Ghz Intel Core 2 Duo with 4GB of RAM.

The aim of our computational analysis is twofold. First, we investigate the behaviour of the optimality gap for all models presented when the number of segments adopted for the piecewise linear approximation of the loss function increases. The term optimality gap is used here to denote the difference between the upper and lower bounds for the expected total cost obtained via Edmundson-Madanski and Jensen’s bounds, respectively. Second, we investigate the computational efficiency of our models and how the number of segments adopted in the piecewise linear approximation of the loss function impacts solution times.

6.1 Test bed

We consider a test bed comprising 810 instances. More specifically, we carried out a full factorial analysis under the following factors. We considered ten different demand patterns illustrated in Fig. 5. The patterns include two life cycle patterns (LCY1 and LCY2), two sinusoidal patterns (SIN1 and SIN2), stationary (STA) and random (RAND) patterns, and four empirical patterns derived from demand data in [14] (EMP1, ..., EMP4). Fixed ordering cost a takes values in $\{500, 1000, 2000\}$; while proportional unit cost v takes values in $\{2, 5, 10\}$. For models under service level constraints, we let the prescribed service level range in $\{0.8, 0.9, 0.95\}$; for model under a penalty cost scheme we let the penalty cost range in $\{2, 5, 10\}$.

6.2 Backorders

In this section, we concentrate on models presented in Section 4.1. Recall that in these models, when a stockout occurs, all demand is backordered and filled as soon as an adequate supply arrives. We first investigate the behaviour of the optimality gap and of the solution time for normally distributed demand. We then extend the analysis to the case in which demand in different periods follow different probability distributions.

6.2.1 Normal distribution

In this section, we assume that demand d_t in each period t to be normally distributed with mean \tilde{d}_t and standard deviation σ_{d_t} , and let the coefficient of variation $c = \tilde{d}_t/\sigma_{d_t}$; we let $c \in \{0.10, 0.20, 0.30\}$. Expected values of the demand in each period are illustrated in Fig. 5 for each of the ten patterns considered. As discussed, when demand is normally distributed, general purpose linearisation parameters can be precomputed and immediately used in our models [24].

In Fig. 6 we report, for each model discussed, boxplots illustrating the optimality gap trend for different number of segments used in the piecewise linear approximation. It should be noted that the y-axis is displayed in logarithmic scale. This shows that the optimality gap shrinks exponentially fast in the number of segments regardless of the model or parameter setting considered.

It is interesting to observe that a number of instances were found infeasible by the MILP model under β^{cyc} service level when two segments for the piecewise approximation were used — note the very large optimality gap. This is due to the fact that with only two segments the approximation error for the β^{cyc} service level was too large and no order-up-to-level could be found to enforce a service level as high as the prescribed one.

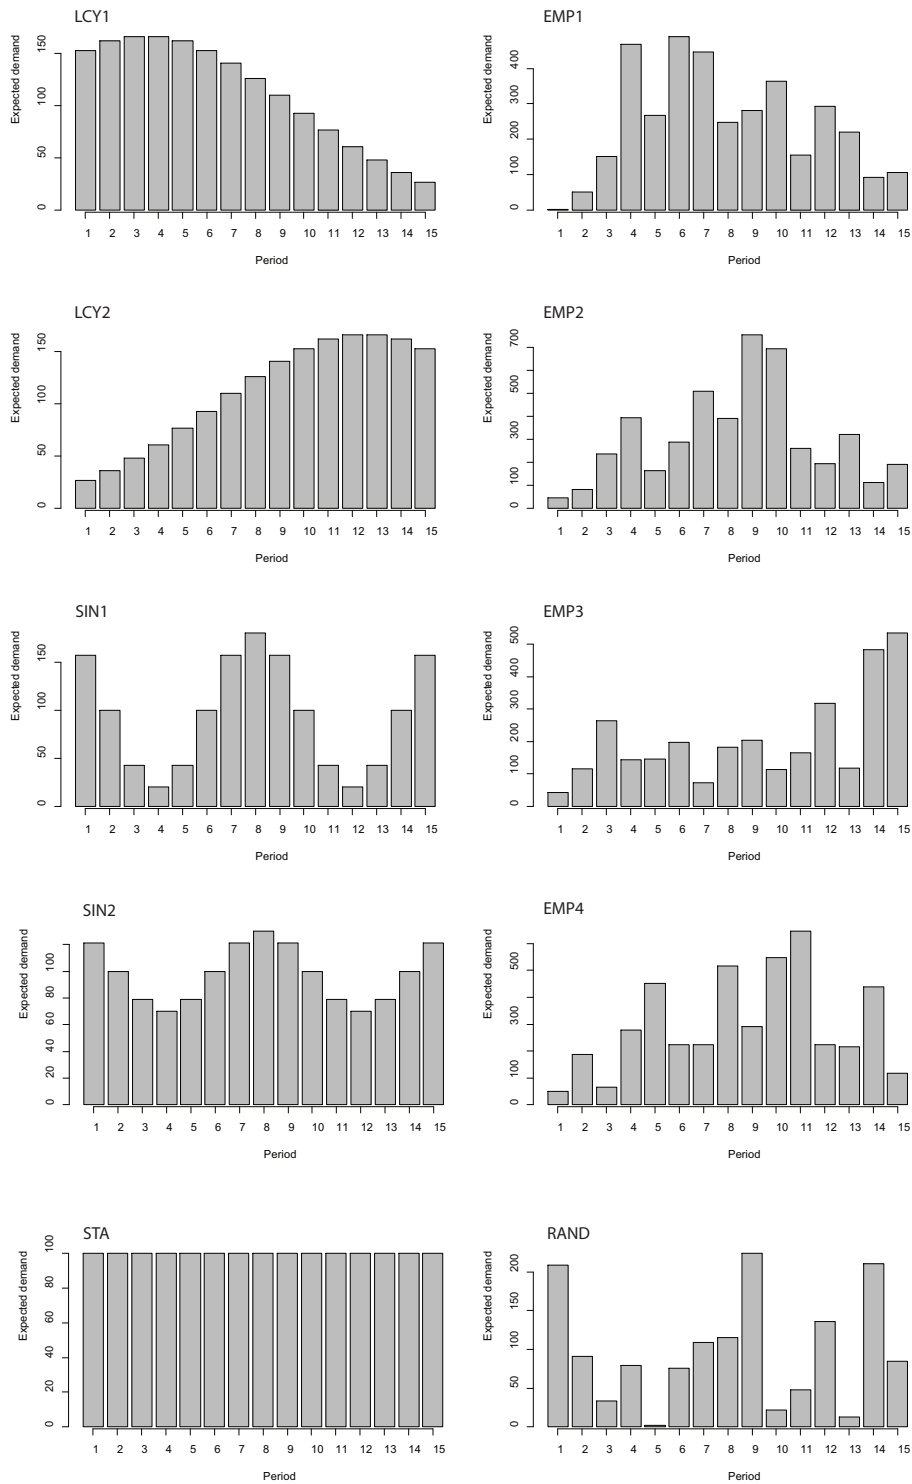


Figure 5: Demand patterns in our computational analysis; the values presented denote the expected demand \tilde{d}_t in each period t of the planning horizon.

Period	Distribution
1,5,9,13	normal with mean \tilde{d}_t and $\sigma_{d_t} = 0.3 \cdot \tilde{d}_t$
2,6,10,14	Poisson with mean \tilde{d}_t
3,7,11,15	exponential with mean \tilde{d}_t
4,8,12	uniformly distributed in $[0, 2\tilde{d}_t]$

Table 1: Probability distribution of the demand in each period of the planning horizon.

In Fig. 7 we report, for each model discussed, boxplots illustrating the computational time trend for different number of segments used in the piecewise linear approximation. Computational times are only slightly affected by the number of segments in the approximation. Furthermore, all instances could be solved in just a few seconds.

6.2.2 Generic distribution

We now extend the analysis to the case in which demand in different periods follow different probability distributions. The test bed remains the same previously analysed. However, the probability distribution of the demand in each period is now detailed in Table 1. To the best of our knowledge, this is the first study in which a numerical analysis on demand that is not normally distributed is presented. We limit our analysis to the model that operates under penalty cost scheme (Section 4.3). We do this for two reasons: firstly, because our local search procedure is computationally quite intensive and a larger test bed would have taken considerable time; secondly, because the model that operates under penalty cost scheme embeds both the first order loss function — employed to compute expected shortages — and the complementary first order loss function — employed to compute holding costs; therefore we expect that results obtained for this model to be sufficiently representative of the overall degree of approximation attained by our piecewise linear approximation of these two functions.

As discussed in Section 3 when demand follows a generic probability distribution we must compute dedicated linearisation parameters for the piecewise first order loss function. In Fig. 8 we consider two possible strategies for computing these parameters. (Fig. 8 - a) illustrates results for the case in which we split the support of the demand ω uniformly into W disjoint compact subregions $\Omega_1, \dots, \Omega_W$ such that $p_i = \Pr\{\omega \in \Omega_i\} = 1/W$. (Fig. 8 - b) illustrates results for the case in which a local search procedure involving a coordinate descent from the most promising partition obtained via simple random sampling is employed to find a good partition into W disjoint compact subregions that minimises the maximum approximation error.




Figure 6: Boxplots illustrating the optimality gap trend for different number of segments used in the piecewise linear approximation.

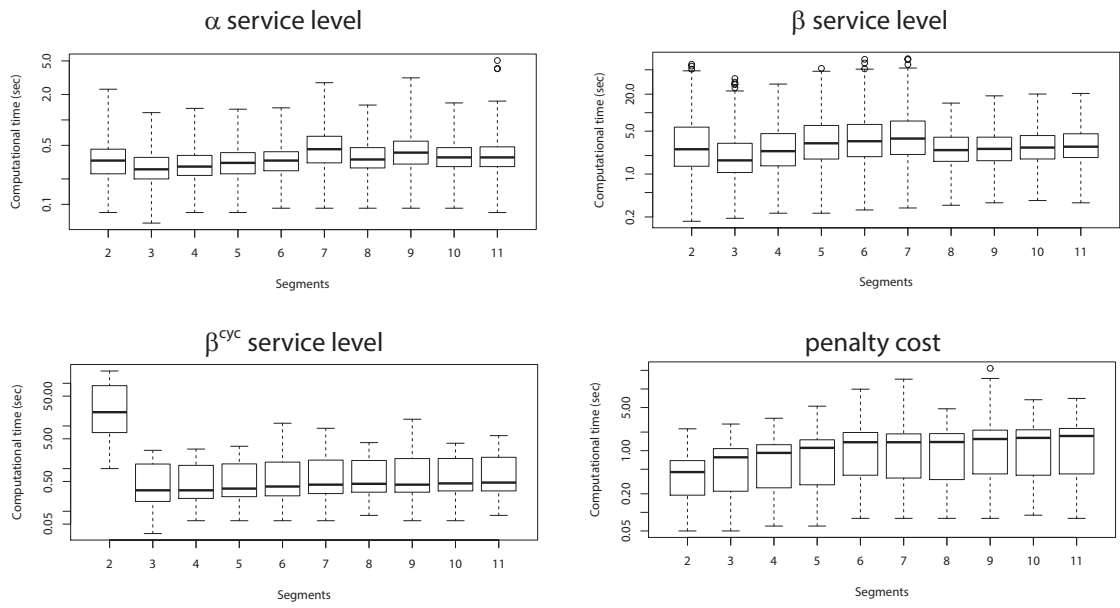


Figure 7: Boxplots illustrating the computational time trend for different number of segments used in the piecewise linear approximation.

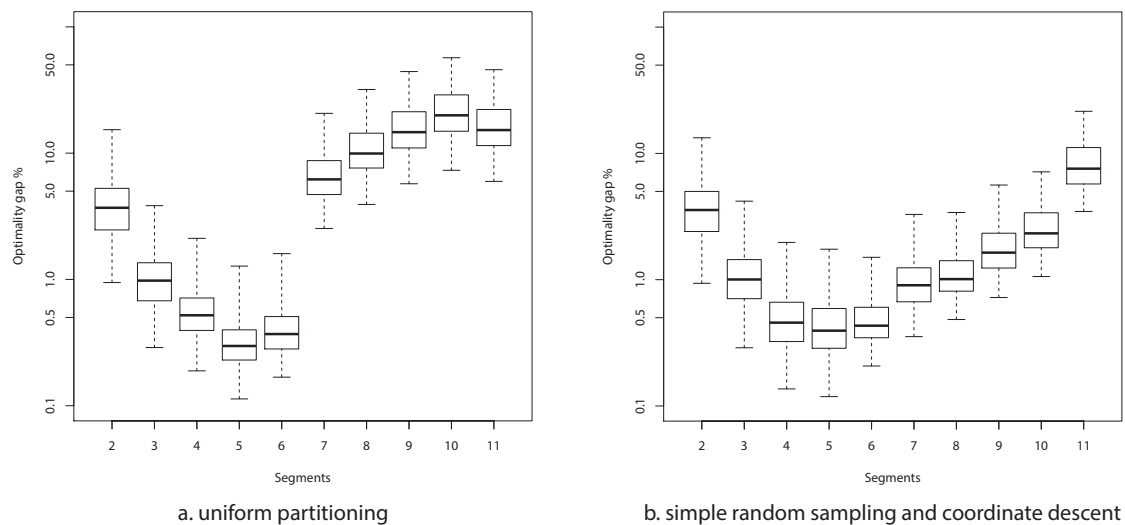


Figure 8: Boxplots illustrating the optimality gap trend for different number of segments used in the piecewise linear approximation.

The results presented reveal that a uniform partitioning of the support of the demand provides acceptable results — optimality gap around 0.5% when the number of segments in the linearisation is low. However, as the number of segments increases, the performance this strategy deteriorates exponentially — note that axis are in logarithmic scale. Local search may help attaining better performances when the number of segments increases. In our specific example, our simple local search strategy, which used a population size $\mathcal{S} = 500W$ for the simple random sampling and a step size $\epsilon = 0.002$ in the coordinate descent, attained up to an order of magnitude improvement in the optimality gap. However, as discussed in [17], this problem is computationally challenging and future research should investigate global optimisation algorithms to compute an optimal partitioning.

A final remark should be made on the computational efficiency of our approach under a generic demand distribution. If one adopts a uniform partitioning of the support then the computational efficiency of the model is essentially identical to the case in which demand is normally distributed. In other words, most of the instances can be solved in few seconds as shown in Fig. 7. If, however, a local search procedure is used to partition the support of the demand then the computational performance of our approach will of course depend on the efficiency of this local search strategy. In our analysis, the local search strategy effectively became the bottleneck, since finding good linearisation parameters for a specific instance took a time that varied from few minutes up to an hour for large W . Also in this case we believe that substantial improvements in computational performances may be achieved via more effective global optimisation algorithms.

6.3 Lost sales

We finally extended our analysis to the case in which demand that occurs when the system is out of stock is lost. Models that operate under this settings were discussed in Section 5. We analyse the case in which demand d_t in each period t is normally distributed with coefficient of variation $c \in \{0.10, 0.20, 0.30\}$.

As shown in Fig. 9 also in this case the optimality gap shrinks exponentially fast for all models considered when the number of segments in the linearisation increases. Furthermore, as shown in Fig. 10 also in this case all instances could be solved in few seconds. We do not discuss models that operate under lost sales and generic probability distributions since results obtained were comparable to those already discussed.

7 Conclusions

We developed MILP formulations for the non-stationary stochastic lot sizing problem. Our formulations exploit a piecewise linearisation of the first order loss function and of its complementary function. We modelled a number of variants of this problem under different service measures: α service level constraints in each period; a penalty cost oriented formulation; a prescribed β^{cyc} service level for each replenishment cycle independently; and a classical β service level, as found in the literature. Our models can be easily adapted to operate in a lost sales setting. Our approach has a number of advantages with respect to other existing approaches in the literature. It is versatile, as it enables seamless modelling of several variants of this problem. It is fully linear and, for the special case in which demand in each period is normally distributed, it does not require an offline evaluation of piecewise linearisation coefficients, as these can be derived from a standard table [see e.g. 24]. As shown in our computational experience, viable linearisation parameters for generic demand distributions can be derived by partitioning the support of the demand uniformly or via local search. Another advantage with respect to other existing approaches is that our models bound from above and below the cost of an optimal plan; by increasing the number of segments in the piecewise linear approximation precision can be improved ad libitum. In our extensive computational experience we demonstrated that optimality gaps shrink exponentially fast in the number of segments used by the piecewise linearisation and that all models developed can be generally solved in a few seconds when up to eleven segments are used in the linearisation.

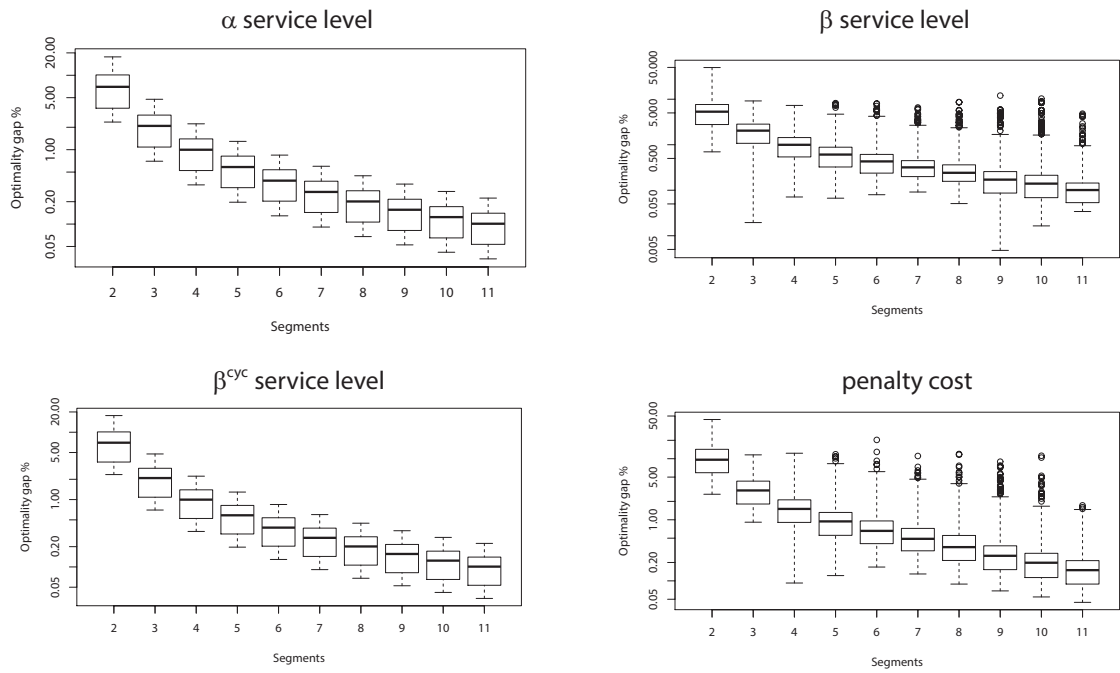


Figure 9: Boxplots illustrating the optimality gap trend for different number of segments used in the piecewise linear approximation.




Figure 10: Boxplots illustrating the computational time trend for different number of segments used in the piecewise linear approximation.

Acknowledgements

S.A. Tarim and O.A. Kilic were supported by Hacettepe University (HU-BAB) and the Scientific and Technological Research Council of Turkey (TUBITAK) under grant No. 110M500. R. Rossi is supported by the University of Edinburgh CHSS Challenge Investment Fund and by the European Community's Seventh Framework Programme (FP7) under grant no 244994 (project VEG-i-TRADE).

References

- [1] Ronald G. Askin. A procedure for production lot sizing with probabilistic dynamic demand. *AIIE Transactions*, 13(2):132–137, June 1981. doi: 10.1080/05695558108974545.
- [2] Sven Axsäter. A simple procedure for determining order quantities under a fill rate constraint and normally distributed lead-time demand. *European Journal of Operational Research*, 174(1):480–491, October 2006. ISSN 03772217. doi: 10.1016/j.ejor.2005.01.037.
- [3] Sven Axsäter. *Inventory Control (International Series in Operations Research & Management Science)*. Springer, 2nd ed. edition, November 2010. ISBN 1441941185.
- [4] Sven Axsäter. When is it feasible to model low discrete demand by a normal distribution? *OR Spectrum*, 35(1):153–162, December 2013. ISSN 0171-6468. doi: 10.1007/s00291-011-0278-8.
- [5] Srinivas Bollapragada and Thomas E. Morton. A simple heuristic for computing non-stationary (s, S) policies. *Operations Research*, 47(4):576–584, 1999.
- [6] James H. Bookbinder and Jin Y. Tan. Strategies for the probabilistic Lot-Sizing problem with Service-Level constraints. *Management Science*, 34(9): 1096–1108, 1988. ISSN 00251909. doi: 10.2307/2632072.
- [7] Jiahua Chen, Dennis K. J. Lin, and Douglas J. Thomas. On the single item fill rate for a finite horizon. *Operations Research Letters*, 31(2):119–123, March 2003. ISSN 01676377. doi: 10.1016/s0167-6377(02)00233-x.
- [8] Momčilo M. Gavrilović. Optimal approximation of convex curves by functions which are piecewise linear. *Journal of Mathematical Analysis and Applications*, 52(2):260–282, November 1975. ISSN 0022247X. doi: 10.1016/0022-247x(75)90095-5.

- [9] George H. Hadley and Thomson M. Whitin. *Analysis of Inventory Systems*. Prentice Hall, 1963. ISBN 0130329533.
- [10] Stefan Helber, Florian Sahling, and Katja Schimmelpfeng. Dynamic capacitated lot sizing with random demand and dynamic safety stocks. *OR Spectrum*, 35(1):75–105, February 2013. ISSN 0171-6468. doi: 10.1007/s00291-012-0283-6.
- [11] Alyson Imamoto and Ben Tang. Optimal piecewise linear approximation of convex functions. In S. I. Ao, Craig Douglas, W. S. Grundfest, Lee Schruben, and Jon Burgstone, editors, *Proceedings of the World Congress on Engineering and Computer Science 2008 WCECS 2008, October 22 - 24, 2008, San Francisco, USA*, pages 1191–1194. IAENG, 2008.
- [12] Peter Kall and Stein W. Wallace. *Stochastic Programming (Wiley Interscience Series in Systems and Optimization)*. John Wiley & Sons, August 1994. ISBN 0471951080.
- [13] Onur A. Kilic and S. Armagan Tarim. An investigation of setup instability in non-stationary stochastic inventory systems. *International Journal of Production Economics*, 133(1):286–292, September 2011. ISSN 09255273. doi: 10.1016/j.ijpe.2010.04.021.
- [14] Abbas A. Kurawarwala and Hirofumi Matsuo. Forecasting and inventory management of short life-cycle products. *Operations Research*, 44(1):131–150, 1996. doi: 10.1287/opre.44.1.131.
- [15] Fatih Mutlu, S[ı]la Çetinkaya, and James H. Bookbinder. An analytical model for computing the optimal time-and-quantity-based policy for consolidated shipments. *IIE Transactions*, 42(5):367–377, February 2010. doi: 10.1080/07408170903462368.
- [16] Ulaş Özen, Mustafa K. Doğru, and S. Armagan Tarim. Static-dynamic uncertainty strategy for a single-item stochastic inventory control problem. *Omega*, 40(3):348–357, June 2012. ISSN 03050483. doi: 10.1016/j.omega.2011.08.002.
- [17] Roberto Rossi and Eligius M. T. Hendrix. Piecewise linearisation of the first order loss function for families of arbitrarily distributed random variables. In Eligius M. T. Hendrix, editor, *Proceedings of MAGO 2014, XII Global Optimization Workshop (GOW) September 1 - 4, 2014, Malaga, Spain*, to appear.
- [18] Roberto Rossi, S. Armagan Tarim, Brahim Hnich, and Steven Prestwich. A global chance-constraint for stochastic inventory systems under service level

- constraints. *Constraints*, 13(4):490–517, December 2008. ISSN 1383-7133. doi: 10.1007/s10601-007-9038-4.
- [19] Roberto Rossi, S. Armagan Tarim, Brahim Hnich, and Steven Prestwich. Computing the non-stationary replenishment cycle inventory policy under stochastic supplier lead-times. *International Journal of Production Economics*, 127(1):180–189, September 2010.
- [20] Roberto Rossi, S. Armagan Tarim, and Ramesh Bollapragada. Constraint-Based local search for inventory control under stochastic demand and lead time. *INFORMS Journal on Computing*, February 2011.
- [21] Roberto Rossi, S. Armagan Tarim, Brahim Hnich, and Steven Prestwich. A state space augmentation algorithm for the replenishment cycle inventory policy. *International Journal of Production Economics*, 133(1):377–384, September 2011. ISSN 09255273. doi: 10.1016/j.ijpe.2010.04.017. URL <http://dx.doi.org/10.1016/j.ijpe.2010.04.017>.
- [22] Roberto Rossi, S. Armagan Tarim, and Onur A. Kilic. A note on tempelmeier’s β -service measure under non-stationary stochastic demand, April 2011. URL <http://arxiv.org/abs/1103.1286>.
- [23] Roberto Rossi, S. Armagan Tarim, Brahim Hnich, and Steven Prestwich. Constraint programming for stochastic inventory systems under shortage cost. *Annals of Operations Research*, 195(1):49–71, July 2012. ISSN 0254-5330. doi: 10.1007/s10479-011-0936-x.
- [24] Roberto Rossi, S. Armagan Tarim, Steven Prestwich, and Brahim Hnich. Piecewise linear approximations of the standard normal first order loss function. *Applied Mathematics and Computation*, to appear. doi: 10.1016/j.amc.2014.01.019. URL <http://arxiv.org/abs/1307.1708>.
- [25] Herbert E. Scarf. Optimality of (s, S) policies in the dynamic inventory problem. In K. J. Arrow, S. Karlin, and P. Suppes, editors, *Mathematical Methods in the Social Sciences*, pages 196–202. Stanford University Press, Stanford, CA, 1960.
- [26] Edward A. Silver. Inventory control under a probabilistic time-varying, demand pattern. *A I I E Transactions*, 10(4):371–379, December 1978. doi: 10.1080/05695557808975228.
- [27] Edward A. Silver and Diane P. Bischak. The exact fill rate in a periodic review base stock system under normally distributed demand. *Omega*, 39(3): 346–349, June 2011. ISSN 03050483. doi: 10.1016/j.omega.2010.08.003.

- [28] Edward A. Silver and Harlan C. Meal. 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, 1973.
- [29] Edward A. Silver, David F. Pyke, and Rein Peterson. *Inventory Management and Production Planning and Scheduling, 3rd Edition*. Wiley, 3 edition, January 1998. ISBN 0471119474.
- [30] Lawrence V. Snyder and Zuo-Jun M. Shen. *Fundamentals of Supply Chain Theory*. Wiley, 1 edition, August 2011. ISBN 0470521309.
- [31] S. Armagan Tarim and Brian G. Kingsman. The stochastic dynamic production/inventory lot-sizing problem with service-level constraints. *International Journal of Production Economics*, 88(1):105–119, March 2004. ISSN 09255273. doi: 10.1016/s0925-5273(03)00182-8.
- [32] S. Armagan Tarim and Brian G. Kingsman. Modelling and computing (Rn,Sn) policies for inventory systems with non-stationary stochastic demand. *European Journal of Operational Research*, 174(1):581–599, October 2006. ISSN 03772217. doi: 10.1016/j.ejor.2005.01.053.
- [33] S. Armagan Tarim and Barbara Smith. Constraint programming for computing non-stationary (R,S) inventory policies. *European Journal of Operational Research*, 189(3):1004–1021, September 2008. ISSN 03772217. doi: 10.1016/j.ejor.2006.11.048.
- [34] S. Armagan Tarim, Brahim Hnich, Roberto Rossi, and Steven D. Prestwich. Cost-Based filtering techniques for stochastic inventory control under service level constraints. *Constraints*, 14(2):137–176, June 2009. ISSN 1383-7133. doi: 10.1007/s10601-007-9039-3.
- [35] S. Armagan Tarim, Mustafa K. Dogru, Ulaş Özen, and Roberto Rossi. An efficient computational method for a stochastic dynamic lot-sizing problem under service-level constraints. *European Journal of Operational Research*, 215(3):563–571, December 2011. ISSN 03772217. doi: 10.1016/j.ejor.2011.06.034.
- [36] Horst Tempelmeier. On the stochastic uncapacitated dynamic single-item lotsizing problem with service level constraints. *European Journal of Operational Research*, 181(1):184–194, August 2007. ISSN 03772217. doi: 10.1016/j.ejor.2006.06.009.

- [37] Horst Tempelmeier. A column generation heuristic for dynamic capacitated lot sizing with random demand under a fill rate constraint. *Omega*, 39(6): 627–633, December 2011. ISSN 03050483. doi: 10.1016/j.omega.2011.01.003.
- [38] Horst Tempelmeier. Stochastic lot sizing problems. In J. MacGregor Smith and Barış Tan, editors, *Handbook of Stochastic Models and Analysis of Manufacturing System Operations*, volume 192 of *International Series in Operations Research & Management Science*, pages 313–344. Springer New York, 2013. doi: 10.1007/978-1-4614-6777-9\10.
- [39] Horst Tempelmeier and Sascha Herpers. ABC_β – a heuristic for dynamic capacitated lot sizing with random demand under a fill rate constraint. *International Journal of Production Research*, 48(17): 5181–5193, August 2009. doi: 10.1080/00207540903179782. URL <http://dx.doi.org/10.1080/00207540903179782>.
- [40] Horst Tempelmeier and Sascha Herpers. Dynamic uncapacitated lot sizing with random demand under a fillrate constraint. *European Journal of Operational Research*, 212(3):497–507, August 2011.
- [41] Douglas J. Thomas. Measuring item fill-rate performance in a finite horizon. *Manufacturing & Service Operations Management*, 7(1):74–80, December 2005. ISSN 1526-5498. doi: 10.1287/msom.1040.0064.
- [42] Huseyin Tunc, Onur A. Kilic, S. Armagan Tarim, and Burak Eksioglu. A simple approach for assessing the cost of system nervousness. *International Journal of Production Economics*, 141(2):619–625, February 2013. ISSN 09255273. doi: 10.1016/j.ijpe.2012.09.022.
- [43] Harvey M. Wagner and Thomson M. Whitin. Dynamic version of the economic lot size model. *Management Science*, 5(1):pp. 89–96, 1958. ISSN 00251909.

Appendix

In Appendix I, we present a list of all symbols used in the paper; in Appendix II, we briefly discuss the two approximate methods adopted in this work for computing good linearisation parameters; in Appendix III, we provide an overview of the MILP models discussed; in Appendix IV we discuss special cases that arise when demand is normally distributed.

Appendix I

x	a scalar value
ω	a random variable
Ω	support of ω
$g_\omega(x)$	probability density function of ω , where $x \in \Omega$
$\tilde{\omega}$	expected value of ω
σ_ω	standard deviation of ω
Z	a standard normal random variable
$\phi(x)$	standard normal probability density function
$\Phi(x)$	standard normal cumulative probability distribution function
W	number of regions in a partition of Ω
i	region index ranging in $1, \dots, W$
Ω_i	a compact region of Ω
p_i	$\Pr\{\omega \in \Omega_i\}$
$E[\omega \Omega_i]$	conditional expectation of ω in Ω_i
$\mathcal{L}(x, \omega)$	first order loss function
$\widehat{\mathcal{L}}(x, \omega)$	complementary first order loss function
$\widehat{\mathcal{L}}_{lb}(x, \omega)$	piecewise linear lower bound of $\widehat{\mathcal{L}}(x, \omega)$ with $W + 1$ segments
$\widehat{\mathcal{L}}_{ub}(x, \omega)$	piecewise linear upper bound of $\widehat{\mathcal{L}}(x, \omega)$ with $W + 1$ segments
e_W	maximum approximation error for $\widehat{\mathcal{L}}_{lb}(x, \omega)$ and $\widehat{\mathcal{L}}_{ub}(x, \omega)$
N	periods in the planning horizon
t	period index ranging in $1, \dots, N$
d_t	demand in period t
c	demand coefficient of variation \tilde{d}_t/σ_{d_t}

$g_t(\cdot)$	probability density function of d_t
$G_t(\cdot)$	cumulative distribution function of d_t
a	fixed ordering cost, \$ a per order
v	proportional unit cost, \$ v per unit ordered
s	selling price, \$ s per item
m	margin, \$ $s - v$ per item
h	holding cost, \$ h per unit of inventory carried to the next period
$E[TC]$	expected total cost
I_0	initial inventory, a scalar
I_t	inventory level at the end of period t
Q_t	order quantity issued (and received) at the beginning of period t
δ_t	binary variable set to one if and only if $Q_t > 0$.
α	non stockout probability
β	prescribed fill rate
β^{cycle}	prescribed cycle fill rate
b	penalty cost \$ b per unit short per unit time
M	a very large number
P_{jt}	a binary variable which is set to one if and only if the most recent inventory review before period t was carried out in period j
S_t	order-up-to-level in period t
d_{jt}	random variable representing the convolution $d_j + \dots + d_t$
$G_{d_{jt}}^{-1}(\alpha)$	α -quantile of the inverse cumulative distribution function of d_{jt}
e_W^{jt}	maximum approximation error associated with the piecewise linearisation of $\widehat{\mathcal{L}}(x, d_{jt})$
\tilde{I}_t^{lb}	a lower bound to the true value of $E[\max(I_t, 0)]$
\tilde{I}_t^{ub}	an upper bound to the true value of $E[\max(I_t, 0)]$
\tilde{B}_t^{lb}	a lower bound to the true value of $E[-\min(I_t, 0)]$
\tilde{B}_t^{ub}	an upper bound to the true value of $E[-\min(I_t, 0)]$
\tilde{C}_t^{lb}	a lower bound to the expected total backorders within the replenishment cycle that ends at period t , if any exists. \tilde{C}_t^{lb} is equal to \tilde{B}_t^{lb} if t is the last period of a replenishment cycle
\tilde{C}_t^{ub}	an upper bound to the expected total backorders within the replenishment cycle that ends at period t , if any exists. \tilde{C}_t^{ub} is equal to \tilde{B}_t^{ub} if t is the last period of a replenishment cycle

Appendix II

We briefly discuss implementation details of the two approximate methods employed in this paper to determine linearisation parameters of the complementary first order loss function for a set of random variables.

The first strategy is quite simple. Consider a set of generic random variables $\omega_1, \dots, \omega_N$ with complementary first order loss functions $\widehat{\mathcal{L}}(x, \omega_1), \dots, \widehat{\mathcal{L}}(x, \omega_N)$. From (2) it is clear that, once all p_i are fixed, all $E[\omega|\Omega_k]$ are uniquely determined for each loss function $\widehat{\mathcal{L}}(x, \omega_n)$; this computation can be carried out numerically by using off-the-shelf packages such as Mathematica⁵ or libraries such as SSJ.⁶ If our aim is to derive a piecewise linearisation with $W + 1$ segments, then a possible strategy is to partition the support of each random variable ω_n into W disjoint compact subregions $\Omega_1, \dots, \Omega_W$, such that $\Pr\{\omega \in \Omega_i\} = 1/W$. Once more, these regions can be determined numerically with one of the aforementioned off-the-shelf packages. By recalling that the complementary first order loss function is convex, it follows that the maximum approximation error will be attained at one of the breakpoints of its piecewise linear approximation. Since there are W breakpoints the maximum approximation error associated with the above partitioning can be easily determined for each $\widehat{\mathcal{L}}(x, \omega_1), \dots, \widehat{\mathcal{L}}(x, \omega_N)$.

The second strategy is still simple, but slightly more refined than the first one. A pseudocode for this strategy is presented in Fig. 3. Instead of simply partitioning the support of random variables $\omega_1, \dots, \omega_N$ uniformly into compact subregions associated with the same probability mass, we now first generate random partitions via simple random sampling and then we try to improve the best of these partitioning via a coordinate descent approach. More specifically, by observing that a partitioning of the support of $\omega_1, \dots, \omega_N$ is uniquely determined by a tuple of W probabilities $\langle p_1, \dots, p_W \rangle$ such that $\sum_{i=1}^W p_i = 1$, we can generate \mathcal{S} such tuples (line 1) and then pick tuple $\langle p_1^*, \dots, p_W^* \rangle$ associated with the minimum maximum approximation error over the set of complementary first order loss functions under scrutiny (line 2). Finally, we try to reduce this error by performing local moves involving each p_i^* separately, for $i = 1, \dots, W - 1$ until no further improvement is possible (line 3). Note that if we decrement/increment p_i^* , since $\sum_{i=1}^W p_i = 1$, then one of the other probabilities must be modified accordingly, this explains why the loops ends at $W - 1$; in other words, we have $W - 1$ degrees of freedom in the search.

⁵<http://www.wolfram.com/mathematica/>

⁶<http://www.iro.umontreal.ca/~simardr/ssj/>

Algorithm 1: Coordinate descent to determine a good partitioning of the random variable supports

Data: random variables $\omega_1, \dots, \omega_N$; size \mathcal{S} of the population for simple random sampling; step size ϵ in the coordinate descent.

Result: a tuple $\langle p_1^*, \dots, p_W^* \rangle$
generate \mathcal{S} tuples $\langle p_1^k, \dots, p_W^k \rangle$;

for $k = 1$ **to** \mathcal{S} **do**

- 1 | determine the minimum maximum approximation error over
| $\widehat{\mathcal{L}}(x, \omega_1), \dots, \widehat{\mathcal{L}}(x, \omega_N)$ under partitioning $\langle p_1^k, \dots, p_W^k \rangle$;
- 2 | **if** partitioning $\langle p_1^k, \dots, p_W^k \rangle$ ensures the best error so far **then**
| | record $\langle p_1^k, \dots, p_W^k \rangle$ as $\langle p_1^*, \dots, p_W^* \rangle$;

3 repeat

for $i = 1$ **to** $W - 1$ **do**

- | determine which partitioning between $\langle p_1^*, \dots, p_i^* + \epsilon, \dots, p_W^* - \epsilon \rangle$
 | and $\langle p_1^*, \dots, p_i^* - \epsilon, \dots, p_W^* + \epsilon \rangle$ leads to the best improvement for the
 | minimum maximum approximation error over $\widehat{\mathcal{L}}(x, \omega_1), \dots, \widehat{\mathcal{L}}(x, \omega_N)$;
- | update $\langle p_1^*, \dots, p_i^*, \dots, p_W^* \rangle$ to reflect the best current partitioning;

until the minimum maximum approximation error cannot be improved;

Appendix III

An overview of the MILP models discussed; numbers refer to the respective equations in the text.

		Lower bound							Upper bound								
		Objective	Subject to							Objective	Subject to						
Backorders	α service level	(20)	(12) (13) (14) (15) (16) (17) (18)	(21)	(12) (13) (14) (15) (16) (17) (18)												
	penalty cost	(24)	(12) (13) (15) (16) (17) (18) (22) (26)	(25)	(12) (13) (15) (16) (17) (18) (23) (27)												
	β^{cyc} service level	(20)	(12) (13) (15) (16) (17) (18) (22) (26) (29)	(21)	(12) (13) (15) (16) (17) (18) (23) (27) (30)												
	β service level	(20)	(12) (13) (15) (16) (17) (18) (22) (26) (33) (35)	(21)	(12) (13) (15) (16) (17) (18) (23) (27) (34) (36)												
Lost sales	α service level	(43)	(22) (44) (45) (46) (47) (48) (49) (50) (51) (52) (53)		replace \tilde{I}_t^{ub} , (22) with \tilde{I}_t^{lb} , (23)												
	penalty cost	(54)	(22) (26) (44) (45) (46) (47) (48) (50) (51) (52) (53)		replace \tilde{I}_t^{ub} , (22), (26), with \tilde{I}_t^{lb} , (23), (27)												
	β^{cyc} service level	(43)	(22) (26) (29) (44) (45) (46) (47) (48) (50) (51) (52) (53)		replace I_t^{ub} , (22), (26), (29) with \tilde{I}_t^{lb} , (23), (27), (30)												
	β service level	(43)	(22) (26) (33) (35) (44) (45) (46) (47) (48) (50) (51) (52) (53)		replace I_t^{ub} , (22), (26), (33), (35) with \tilde{I}_t^{lb} , (23), (27), (34), (36)												

Appendix IV

In this appendix we discuss model variants for the case in which demand in each period is normally distributed. The assumption that demand is normally distributed plays a prominent role in inventory theory [see e.g. 4] and most of existing works on stochastic lot sizing focus on this distribution. An important property of the first order normal loss function is that it can be derived from its standard counterpart. Let ω be a normally distributed random variable with mean μ and standard deviation σ . Let $\phi(x)$ be the standard normal probability density function and $\Phi(x)$ the respective cumulative distribution function.

Lemma 4. *The complementary first order loss function of ω can be expressed in terms of the standard normal cumulative distribution function as*

$$\widehat{\mathcal{L}}(x, \omega) = \sigma \int_{-\infty}^{\frac{x-\mu}{\sigma}} \Phi(t) dt = \sigma \widehat{\mathcal{L}}\left(\frac{x-\mu}{\sigma}, Z\right), \quad (55)$$

where Z is a standard normal random variable.

Rossi et al. [24] discussed how to obtain an optimal partitioning of the support under a framework that minimises the maximum approximation error. The same work reports standard linearisation parameters for the case in which ω is a standard normal random variable.

α service level constraints

If demand in each period is normally distributed, we can exploit Lemma 4 to reduce the number of linearisation parameters that are needed in the model. Consider a partition of the support Ω of a standard normal random variable Z into W adjacent regions Ω_i . Recall that $p_i = \Pr\{Z \in \Omega_i\}$, by exploiting the Jensen's piecewise linear lower bound, we introduce the following constraints in the model

$$\tilde{I}_t^{lb} \geq \tilde{I}_t \sum_{k=1}^i p_k - \sum_{j=1}^t \left(\sum_{k=1}^i p_k \mathbb{E}[Z|\Omega_i] \right) P_{jt} \sigma_{d_{jt}} \quad t = 1, \dots, N; \quad i = 1, \dots, W$$

where $\sigma_{d_{jt}}$ denotes the standard deviation of $d_j + \dots + d_t$ and $\tilde{I}_t^{lb} \geq 0$. This expression follows immediately from Lemma 2 and Lemma 4: consider a replenishment cycle covering periods j, \dots, t and associated order-up-to-level S . We aim to enforce $\tilde{I}_t^{lb} \geq \sigma \widehat{\mathcal{L}}_{lb}^i \left((S - \tilde{d}_{jt}) / \sigma_{d_{jt}}, Z \right)$ for all $i = 1, \dots, W$. Observe that $S - \tilde{d}_{jt} = \tilde{I}_t$,

the above expression follows immediately. We then derive \tilde{I}_t^{ub} from \tilde{I}_t

$$\tilde{I}_t^{ub} \geq \tilde{I}_t \sum_{k=1}^i p_k + \sum_{j=1}^t \left(e_W - \sum_{k=1}^i p_k \mathbb{E}[Z|\Omega_i] \right) P_{jt} \sigma_{d_{jt}} \quad \begin{array}{l} t = 1, \dots, N, \\ i = 1, \dots, W; \end{array}$$

where $\tilde{I}_t^{ub} \geq \sum_{j=1}^t e_W P_{jt} \sigma_{d_{jt}}$ for $t = 1, \dots, N$; and e_W denotes the maximum approximation error associated with a partition comprising W regions; linearisation parameters that minimise the maximum approximation error e_W for a given number $W + 1$ of segments, where $W = 1, \dots, 10$, can be found in [24].

A final remark that is worth making is the fact that it is possible to replace e_W with $e_W/2$ in the above equations to obtain an approximation of the first order loss function that minimises the maximum absolute error. However, this approximation does not allow one to establish if the cost produced by the model is an upper or a lower bound for the true cost of an optimal plan.

Penalty cost scheme

We discuss how to handle the case in which demand is normally distributed by exploiting Lemma 1. We obtain \tilde{B}_t^{lb} and \tilde{B}_t^{ub} from \tilde{I}_t by exploiting the connection between the Jensen's piecewise linear lower bound and the piecewise linear upper bound to the first order loss function (Lemma 1).

$$\tilde{B}_t^{lb} \geq -\tilde{I}_t + \tilde{I}_t \sum_{k=1}^i p_k - \sum_{j=1}^t \left(\sum_{k=1}^i p_k \mathbb{E}[Z|\Omega_i] \right) P_{jt} \sigma_{d_{jt}} \quad \begin{array}{l} t = 1, \dots, N, \\ i = 1, \dots, W; \end{array} \quad (56)$$

where $\tilde{B}_t^{ub} \geq -\tilde{I}_t$ and

$$\tilde{B}_t^{ub} \geq -\tilde{I}_t + \tilde{I}_t \sum_{k=1}^i p_k + \sum_{j=1}^t \left(e_W - \sum_{k=1}^i p_k \mathbb{E}[Z|\Omega_i] \right) P_{jt} \sigma_{d_{jt}} \quad \begin{array}{l} t = 1, \dots, N, \\ i = 1, \dots, W; \end{array} \quad (57)$$

where $\tilde{B}_t^{ub} \geq -\tilde{I}_t + \sum_{j=1}^t e_W P_{jt} \sigma_{d_{jt}}$.