Omega 50 (2015) 126–140

Contents lists available at ScienceDirect

Omega

journal homepage: www.elsevier.com/locate/omega

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

Roberto Rossi [a,*,1], Onur A. Kilic [b,2], S. Armagan Tarim [b,c,2,3]

[a] Business School, University of Edinburgh, 29 Buccleuch Place, EH8 9JS Edinburgh, UK
[b] Institute of Population Studies, Hacettepe University, Turkey
[c] Insight Centre for Data Analytics, University College Cork, Ireland

ARTICLE INFO

Article history:
Received 5 February 2014
Accepted 7 August 2014
This manuscript was processed by Associate Editor Kuhn
Available online 20 August 2014

Keywords:
Stochastic lot sizing
Static–dynamic uncertainty
First order loss function
Non-stockout probability
Fill rate
Penalty cost
Piecewise linearisation

ABSTRACT

In this paper, we develop a unified mixed integer linear modelling approach to compute near-optimal policy parameters for the non-stationary stochastic lot sizing problem under static–dynamic uncertainty strategy. The proposed approach applies to settings in which unmet demand is backordered or lost; and it can accommodate variants of the problem for which the quality of service is captured by means of backorder penalty costs, non-stockout probabilities, or fill rate constraints. This approach has a number of advantages with respect to existing methods in the literature: it enables seamless modelling of different variants of the stochastic lot sizing problem, some of which have been previously tackled via ad hoc solution methods and some others that have not yet been addressed in the literature; and it produces an accurate estimation of the expected total cost, expressed in terms of upper and lower bounds based on piecewise linearisation of the first order loss function. We illustrate the effectiveness and flexibility of the proposed approach by means of a computational study.

1. Introduction

We consider the non-stationary stochastic lot sizing problem — the stochastic extension of the well-known dynamic lot sizing problem [1]. 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 [2] 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, since 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 [3]. The static uncertainty is appealing in material requirement planning systems, for which order synchronisation is a key concern [4]. The static–dynamic uncertainty strategy has advantages in organising joint replenishments and shipment consolidation [5–8].

In this study, we focus our attention on the static–dynamic uncertainty strategy, which offers a stable replenishment plan while effectively hedging against uncertainty [9,10]. 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. [11] 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

☆This manuscript was processed by Associate Editor Kuhn.

* Corresponding author. Tel.: +44 131 6515239; fax: +44 131 650 8077.
E-mail addresses: roberto.rossi@ed.ac.uk (R. Rossi), onuralp@hacettepe.edu.tr (O.A. Kilic), armagan.tarim@hacettepe.edu.tr (S.A. Tarim).

[1] The author is supported by the University of Edinburgh CHSS Challenge Investment Fund.

[2] The authors were partially supported by the Scientific and Technological Research Council of Turkey (TUBITAK) under Grant no. 110M500.

[3] This publication has emanated from research supported in part by a research grant from Science Foundation Ireland (SFI) under Grant number SFI/12/RC/2289.

http://dx.doi.org/10.1016/j.omega.2014.08.003
0305-0483/

review period — these decisions must be fixed at the beginning of the planning horizon.

We introduce a unified modelling approach that captures several variants of the problem and that is based on standard mixed-integer linear programming (MILP) 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. [5, pp. 244–246]): penalty cost per unit short per period, non-stockout probability[4] (α service level), cycle fill rate[5] (β^{cyc} service level), and fill rate[6] (β service level). Our models build on recently introduced piecewise linear upper and lower bounds for the first order loss function and its complementary function [12], which are based on distribution independent bounding techniques from stochastic programming: Jensen's and Edmundson–Madanski's inequalities [13, pp. 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 [12].

Our contributions to the inventory control literature are the following:

- we develop a unified MILP modelling approach that enables seamless modelling of the non-stationary stochastic lot sizing problem under each of the four measures of service quality discussed;
- 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 Hadley and Whitin [14], Silver et al. [5], and Axsäter [15].
- 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;
- 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 present an extensive computational study to show that (i) whenever other state-of-the-art approximations exist, our models feature comparable optimality gaps; (ii) the linearisation gap — the term linearisation 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 — shrinks exponentially fast as the number of segments in the piecewise linearisation increases; and (iii) the number of segments adopted only marginally affects computational efficiency.

[4] A lower bound on the non-stockout probability in any period over the planning horizon.

[5] A lower bound on the expected fraction of demand that is routinely satisfied from stock for each replenishment cycle, i.e. the time interval between two successive inventory reviews.

[6] A lower bound on the expected fraction of demand that is routinely satisfied from stock over the planning horizon.

2. Literature survey

Due to its practical relevance, a large body of literature has emerged on the static–dynamic uncertainty strategy over the last few decades. Here, we review some key studies which are of particular importance in the context of our work, and reflect upon our contribution. All the models discussed in the following sections operate under a static–dynamic uncertainty strategy. To keep our discussion focused, in what follows we only survey studies related to the static–dynamic uncertainty strategy and we disregard those addressing the static uncertainty (see e.g. [16–19]) or the dynamic uncertainty strategy (see e.g. [20,21]).

Early works on the stochastic lot sizing problem concentrated on easy-to-compute heuristics. Silver [22] and Askin [23] 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 [24].

Bookbinder and Tan [2] 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 modelling 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 [25] discussed the first MILP formulation under α service level constraints. In contrast to [2], this formulation simultaneously determines the replenishment schedule and corresponding order-up-to levels. Efficient reformulations operating under the same assumptions were discussed in [26–28]. Rossi et al. [27] developed a state space augmentation approach; Tarim et al. [26] implemented a branch and bound algorithm; and Tunc et al. [28] developed an effective MILP reformulation. In addition, Constraint Programming reformulations based on a novel modelling tool, i.e. global chance constraints, were discussed in [29,30]. Finally, an exact, although computationally intensive, Constraint Programming approach was discussed in [31]. Extensions to the case of a stochastic delivery lead time were discussed in [32,33].

Tarim and Kingsman [34] developed the first MILP formulation for the case in which service quality is modelled using a penalty cost scheme. Rossi et al. [35] discussed an efficient Constraint Programming reformulation exploiting optimization oriented global stochastic constraints.

Özen et al. [11] 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” heuristic 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