University of Groningen

The benefit of receding horizon control

Dural-Selcuk, Gozdem; Rossi, Roberto; Kilic, Onur A.; Tarim, S. Armagan

Published in:
Omega: The International Journal of Management Science

DOI:
[10.1016/j.omega.2019.07.007](https://doi.org/10.1016/j.omega.2019.07.007)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2020

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Dural-Selcuk, G., Rossi, R., Kilic, O. A., & Tarim, S. A. (2020). The benefit of receding horizon control: Near-optimal policies for stochastic inventory control. *Omega: The International Journal of Management Science*, 97, [102091]. <https://doi.org/10.1016/j.omega.2019.07.007>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

The benefit of receding horizon control: Near-optimal policies for stochastic inventory control[☆]

Gozdem Dural-Selcuk^a, Roberto Rossi^{b,*}, Onur A. Kilic^c, S. Armagan Tarim^d

^aIndustrial Engineering, Atılım University, Ankara, Turkey

^bBusiness School, University of Edinburgh, Edinburgh, UK

^cFaculty of Economics and Business, University of Groningen, Groningen, the Netherlands

^dCork University Business School, University College Cork, Ireland

ARTICLE INFO

Article history:

Received 24 February 2019

Accepted 23 July 2019

Available online 24 July 2019

Keywords:

Stochastic lot sizing

Static uncertainty

Dynamic uncertainty

Static-dynamic uncertainty

Receding horizon control

ABSTRACT

In this paper we address the single-item, single-stocking point, non-stationary stochastic lot-sizing problem under backorder costs. It is well known that the (s, S) policy provides the optimal control for such inventory systems. However the computational difficulties and the nervousness inherent in (s, S) paved the way for the development of various near-optimal inventory control policies. We provide a systematic comparison of these policies and present their expected cost performances. We further show that when these policies are used in a receding horizon framework the cost performances improve considerably and differences among policies become insignificant.

© 2019 Elsevier Ltd. All rights reserved.

1. Introduction

The aim of lot sizing is to determine replenishment quantities and their timings to minimise total inventory cost while balancing supply and demand [1]. A major stream within the lot-sizing literature addresses demand uncertainty in a dynamic environment, i.e. non-stationary stochastic demand. The stochastic nature of demand is associated with market uncertainty that could emerge from several factors, such as pricing, advertising, channel development, manufacturing and inventory management [2]. It may also originate from the interdependency of end customer's demand with supply yield [3,4]. The shortening of product life cycles, on the other hand, accounts for the non-stationarity feature of demand: as product life cycles get shorter, demand rates are subject to quick changes over time. Furthermore, seasonality and trend may cause demand rates to change within the phases of the product life cycle. As a result of being stochastic and dynamic, demand usually has varying probability distributions over time [5]. This fact implies that market uncertainty necessitates inventory control policies with time-varying parameters. As shown in [6], it is costly to adopt stationary policies under non-stationary demand.

Our study focuses on the single-item, single-stocking point, non-stationary stochastic lot-sizing problem under backorder costs. This is a well-studied problem in the literature. In his seminal work [7], Scarf demonstrated the optimality of (s, S) -type policies under this problem setting. The computation of optimal stationary (s, S) policy parameters was investigated shortly thereafter in [8]. However, although the form of the optimal policy has been known for a long time, computing optimal non-stationary (s, S) policy parameters still remains a computationally challenging task. Moreover, the instability in order plans (i.e. nervousness) associated with this control policy is also a key concern [9].

Early heuristics that were proposed for this problem setting include [10], which represents an extension of the classical Silver-Meal [11] algorithm to a stochastic setting, [12,13]. The framework introduced in [12], which comprises three types of policies, static-uncertainty, static-dynamic uncertainty, and dynamic uncertainty, motivated a number of follow up works, such as [14–18]. The former studies focus on the static-uncertainty policy; whereas the later ones focus on the static-dynamic uncertainty strategy.

This work considers only backorder costs related to shortages. Other variants of the problem investigated in the literature considered α service level constraints [19–21] and β (i.e. fill rate) service level constraints [22]. These variants are left out of the scope of this work since benchmark optimal policies are not available in the literature. In an early work [23], Bookbinder and H'ng compared the performance of six separate lot-sizing strategies for use in a rolling schedule. To the best of our knowledge, there is no similar

[☆] This manuscript was processed by Associate Editor B. Keskin.

* Corresponding author.

E-mail addresses: gozdemduralselecuk@gmail.com (G. Dural-Selcuk), roberto.rossi@ed.ac.uk (R. Rossi), o.a.kilic@rug.nl (O.A. Kilic), armagan.tarim@ucc.ie (S.A. Tarim).

study focusing on stochastic lot sizing heuristics under backorder cost. The current manuscript aims at filling this gap by making the following contributions to the literature:

- We present a systematic comparison on expected cost performances of three lot-sizing policies for the single-item, single-stocking point, non-stationary stochastic lot-sizing problem under backorder costs.
- We compare the cost performance of these policies in a receding horizon deployment [24], which is widely used in the lot-sizing literature (i.e. [25,26]). Receding horizon control proceeds as follows: a replenishment plan is obtained for the entire planning horizon, but only the imminent replenishment decision is implemented, and a re-planning is done at the beginning of each period for the rest of the planning horizon.
- We derive the following managerial insights: we show to what extent the receding horizon scheme improves the cost performances of various policies, and how the cost performance of different policies become insignificant when they are implemented in a receding horizon framework.

The rest of the paper is organised as follows. Section 2 gives a formal definition of the problem under consideration. Section 3 gives a detailed overview of the existing approaches in the literature. Section 4 presents the numerical study and the results obtained. Finally, Section 5 concludes with final remarks.

2. Problem definition

We focus on the single-item, single-stocking location, stochastic non-stationary lot sizing problem; which is defined as follows. The planning horizon is finite and it consists of N discrete time periods. Individual period demands $d_1, d_2, \dots, d_N$ are independent random variables with known probability distribution functions not necessarily identically distributed. A holding cost, h is incurred on any unit carried in inventory from one period to the next. Any demand that cannot be satisfied immediately is backordered and satisfied whenever a sufficiently sized replenishment order arrives. A penalty cost b is incurred for each unit of demand backordered per period. A fixed ordering cost K is incurred every time an order is placed.

A mathematical formulation for the problem may be written as follows (see e.g. [14,16]):

$$\min \sum_{n=1}^N \mathbb{E} \{ Kz_n + hx_n^+ + bx_n^- \} \quad (1)$$

subject to

$$x_n = x_{n-1} + q_n - d_n \quad \forall n \in [1, N] \quad (2)$$

$$z_n = \begin{cases} 1 & \text{if } q_n > 0 \\ 0 & \text{otherwise} \end{cases} \quad \forall n \in [1, N] \quad (3)$$

$$z_n \in \{0, 1\}, \quad q_n \in \mathbb{R}_+ \quad \forall n \in [1, N] \quad (4)$$

Here z_n and q_n are the decision variables indicating the replenishment action and the replenishment quantity for period n ; where x_n stands for the end of period inventory position. Also, $x^+ = \max(0, x)$ and $x^- = \max(0, -x)$ respectively stand for the amount of excess stocks and backorders. We assume that inventory levels are accurate and do not include any source of uncertainty arising from information inaccuracy (see e.g. [27,28]).

2.1. Control policies

In [12], it is argued that the aforementioned model leads to different control strategies depending on when replenishment decisions z_n and q_n are made. They discuss three classes of strategies:

dynamic uncertainty, static uncertainty, and static-dynamic uncertainty.

In the *dynamic uncertainty strategy*, the decision maker observes the current inventory position at the beginning of each time period n , i.e. x_{n-1} , and decides whether to place an order, and if so how much to order. Therefore, z_n and q_n are determined at period n .

In the *static uncertainty strategy*, timing and quantity of orders are fixed once and for all at the beginning of the planning horizon. Thus, all q_n and z_n are determined at the outset.

In the *static-dynamic uncertainty strategy*, the complete replenishment schedule is determined at the beginning of the planning horizon, whereas order quantities are decided at the time of replenishments. Hence, all z_n are determined at the outset but the decision on q_n is postponed until period n .

Intuitively, the cost-effectiveness of an inventory control policy is positively correlated with the amount of demand information used at the time of making replenishment decisions. Thus, the dynamic uncertainty strategy has a superior cost performance as compared to static-dynamic and static uncertainty strategies, since it makes replenishment decisions only after observing the actual inventory level.

The cost performance of these strategies is numerically analysed in [9]. The authors show that the static-dynamic uncertainty strategy is very competitive, whereas the static uncertainty performs rather poorly in comparison to the dynamic uncertainty strategy. As we will show in our numerical study, results are substantially different under a receding horizon deployment of these policies.

3. Survey of exact and heuristics approaches in the literature

We review both optimal and heuristic approaches proposed in the literature to compute the parameters of dynamic, static and static-dynamic uncertainty strategies for the stochastic lot sizing problem, with and without receding horizon approach. For the purposes of this study, we specifically address those studies analysing non-stationary stochastic inventory systems operating under a penalty cost scheme for backordered demand.

3.1. Dynamic uncertainty strategy

The literature on dynamic demand dates back to Wagner and Whitin's well-known deterministic dynamic lot sizing model [29]. In his seminal paper [7], Scarf characterizes the optimal control policy for the lot sizing problem under stochastic demand; this characterization holds for stationary as well as non-stationary demand. The optimal policy determines two critical parameters for each period: the re-order level s and the order up-to level S ; for this reason it is named (s, S) policy. The decision maker observes the inventory position at the beginning of a period, and places an order if it is below the re-order level so as to replenish the inventory up to the order up-to level. If the inventory position is above the reorder level no order is placed. The optimal policy follows the dynamic uncertainty strategy because both the timing and the quantity of orders become known only at replenishment epochs.

In [7], Scarf identified the structure of the optimal policy, but finding optimal parameters has remained a computationally intensive task ever since. That is because one needs to recursively compute a continuous cost function in order to obtain the optimal re-order and order-up-to levels for each and every period within the planning horizon. An alternative approach that can be used to tackle the continuity issue is to use a discrete demand distribution (see e.g. [13]). In this case, it is possible to use a discrete state space dynamic program to obtain the optimal cost function. Nevertheless, the resulting procedure is still complex since there is a

very wide range of possible inventory levels that should be considered. A variety of studies address this issue (see e.g. [8,30–32]). However, most of these studies handle the problem under stationary demand and an infinite planning horizon.

There are a few studies suggesting heuristic methods to compute parameters of the dynamic uncertainty strategy. Askin proposes a heuristic method [10] which adapts Silver and Meal's well-known heuristic designed for the deterministic version of the same problem. Askin's heuristic first determines the cycle length, i.e. the number of periods to be covered, and the order-up-to level for all prospective orders through the planning horizon. The length of a replenishment cycle is selected so as to minimize the corresponding average total cost per period. The re-order level associated with a replenishment cycle, on the other hand, is myopically set to an inventory level which minimizes the costs to be incurred during the imminent replenishment cycle. The heuristic determines the re-order levels by means of a trade-off analysis between expected costs per period in cases of ordering and not ordering. In particular, the re-order level of a particular period is set to an inventory level where the difference between the expected costs of ordering and not ordering equals the fixed ordering cost. Bollapragada and Morton [13] also propose a myopic heuristic to compute the parameters of dynamic uncertainty strategy. This heuristic relies on the idea of approximating the non-stationary problem by a series of stationary problems. The heuristic proceeds as follows. First, by means of the method developed by Zheng and Federgruen [32], for all possible values of mean demand, the optimal parameters of the associated stationary problem as well as the expected time between two consecutive orders are obtained and tabulated. Then, for each period, the order-up-to and re-order levels are set to the corresponding optimal parameters of a specific stationary problem which is chosen in such a way that the cumulative mean demands of the stationary and non-stationary problems over the expected reorder cycle of the stationary problem are equal to each other. The relevant values used in this procedure are read from the table generated in the first step of the heuristic. An important drawback of this method is that it cannot account for end-of-horizon effects due to the underlying stationary approximation. Bollapragada and Morton [13] overcome this issue by replacing the heuristic parameters with the optimal ones for some periods at the end of the planning horizon.

3.2. Static uncertainty strategy

The static uncertainty strategy decides timing and quantity of replenishments at the beginning of the planning horizon. It is customary in the literature to denote the timing as R —where R denotes the length of time, in periods, between two consecutive replenishments—and the quantity as Q ; for this reason this policy is often denoted as the (R, Q) policy. Neither timing nor quantity of replenishments are revised in response to demand realisations over the planning horizon. Therefore in this policy demand uncertainty is only considered “statically” while determining appropriate values for R and Q at the beginning of the planning horizon, before any demand realisation is observed. Despite being unable to exploit additional information obtained from past demand realisations, the static uncertainty strategy has the advantage of providing a completely stable production environment, which is appealing in industrial environments characterized by a low degree of flexibility.

Sox [14] studies this strategy and models the problem as a mixed integer non-linear program. He provides a solution algorithm based on a network formulation of the problem where each arc corresponds to a prospective replenishment cycle. Here, arc costs are calculated by using optimal cumulative replenishment quantities up to and including the corresponding replenishment

cycles. Sox also derives some properties of the optimal solution which he uses to increase the efficiency of the proposed algorithm. This algorithm can be regarded as a stochastic extension of the Wagner and Whitin's algorithm [29] with additional feasibility constraints. Vargas employs a similar approach [15] and shows that optimal replenishment quantities for a given replenishment schedule follow a critical ratio rule. Also, he explicitly addresses the special case where demands are normally distributed, and provides a simple, yet very efficient solution algorithm by exploiting the properties of the normal distribution.

3.3. Static-dynamic uncertainty strategy

The static-dynamic uncertainty strategy provides a stable replenishment pattern by fixing the timing of future orders in advance [16]; however, actual order quantities are decided only after demand during periods that precede a given replenishment have been realised. To do so, the decision maker fixes a so-called “order-up-to-level” for each replenishment, this represents the level up to which inventory must be raised by the current order. It is customary in the literature to denote the order-up-to-level as S ; for this reason this policy is often denoted as the (R, S) policy. This policy is less conservative than the static uncertainty strategy, and it can hedge against uncertainty more efficiently [9,26]. Because it offers a fixed order schedule, the static-dynamic strategy is particularly appealing in material requirement planning, joint replenishment, and shipment consolidation environments (see [1,33–35]).

The contributions in this line of research are as follows. Silver proposes a heuristic [36] which can be regarded as the stochastic version of the heuristic in [11]. This heuristic sequentially determines the timing of the next replenishment period starting from the first period of the planning horizon. Here, the expected cost per period is defined as a function of the number of periods the current order is to cover when the associated order-up-to level is myopically determined so as to minimize the expected costs to be incurred until the next replenishment epoch. The procedure postpones the next replenishment period as long as the expected cost per period is decreasing. Penalty costs for backordered demands are not explicitly mentioned and order quantities are rather determined to ensure a desired service level until the next replenishment epoch. The service level implementation makes this solution approach lie beyond the scope of our numerical analysis, which only considers the non-stationary stochastic lot-sizing problem under a penalty cost scheme.

Tarim and Kingsman [16] provide a mixed integer programming (MIP) formulation of the problem. As opposed to the procedure employed by Silver [36], their method determines order-periods and order-up-to levels simultaneously under the penalty cost assumption. This, however, comes with an additional difficulty which stems from the interdependence between order-up-to levels and costs associated with consecutive replenishment cycles, as well as the non-linear cost function. Tarim and Kingsman [16] overcome these issues by making the additional assumption that demands are normally distributed. This enables them to develop a certainty equivalent mixed integer programming model in which a piecewise linear approximation is used to express the non-linear terms in the objective function. Tarim and Kingsman's formulation [16] does not allow expected order sizes to be negative. Hence, if the inventory level at the beginning of a replenishment cycle happens to exceed the corresponding order-up-to level, then the excess inventory is carried forward. Nevertheless, their method implicitly assumes that such instances are rare events and ignores the cost of carrying excess inventories. As such, their approach does not necessarily provide the optimal solution. Rossi et al. [17] develop a stochastic constraint programming model which is mainly equivalent to the MIP model in [16]. However, instead of employ-

Table 1

An overview of the methods considered in the current study; “MIP” denotes the fact that the approach relies on a MIP model.

Author(s)	Abbrev.	Strategy	Approach	Complexity
Askin [10]	Ask	Dynamic	Heuristic	$O(N)$
Bollapragada and Morton [13]	Bol	Dynamic	Heuristic	$O(N) + O(M)$
Tarim and Kingsman [16]	Tar	Static-Dynamic	Heuristic	MIP
Rossi et al. [17]	Ros	Static-Dynamic	Heuristic	MIP
Sox [14]	Sox	Static	Optimal	$O(N^2)$ (DP)
Vargas [15]	Var	Static	Optimal	MIP

ing a piecewise linear approximation, they use the exact non-linear cost function. They also introduce a cost based filtering method (see e.g. [37,38]) which exploits the convexity of the cost-function for a fixed replenishment schedule. The filtering method dynamically produces bounds on the optimal total cost during the search procedure of the constraint program for fixed values of binary variables indicating the timing of replenishment epochs, and hence, leads to significant improvements in the computational performance of the proposed method. Although the formulation in [17] makes use of the exact non-linear cost function, it may not always yield the optimal policy parameters because – as is the case for Tarim and Kingsman’s formulation [16] – it ignores the cost of carrying excess inventories in cases where the actual inventory level exceeds the order-up-to level. More recently, Rossi et al. [39] extended MIP approach in [16] to generic demand distributions and to a number of service level measures.

3.4. Computational efficiency of existing heuristics

We briefly focus on the computational complexity of the approaches considered in our study (Table 1).

With a linear complexity, Askin’s approach [10] is computationally attractive; the approach processes all N periods in the planning horizon sequentially and utilises a closed form critical fractile expression—similar to a newsvendor solution—to determine the optimal order quantity. Bollapragada and Morton’s approach [13] is also very efficient: the heuristic essentially boils down to a line search procedure over a pre-computed table. However, the precomputation of the table has pseudo-polynomial complexity, since table elements should cover all possible values that the expected demand may take at a given period up to the maximum expected demand M . Three works, [15–17] rely on MIP formulations; while [14] proposed a heuristic with quadratic time complexity that takes the form of a forward dynamic programming (DP) algorithm.

All the approaches, including those based on MIP formulations, are computationally very efficient and can solve realistic instances with up to 30 periods in fractions of a second, as illustrated in the studies listed in Table 1 as well as in related follow up works. Conversely, an exact stochastic dynamic programming approach has pseudo-polynomial complexity and, depending on the demand and state space discretization/truncation adopted, may be computationally very expensive. Since differences in computational performance of all approaches listed in Table 1 are negligible for most practical purposes, in what follows we will only focus on their relative cost performance.

4. Numerical study

This numerical study aims to present a cost based comparison of stochastic lot-sizing strategies under non-stationary stochastic demand and penalty cost assumptions on a common test bed from the literature. The list of the approaches in the scope of this numerical experiment is given in Table 1.

In addition to comparing the approaches as such, we also compare them under a *receding horizon* deployment [24], which is widely used in the lot-sizing literature (i.e. [25,26]). The receding horizon control in inventory problems proceeds as follows: the replenishment plans are made over the entire planning horizon, but only the imminent replenishment decision is implemented and a re-planning is done at the beginning of each period for the rest of the planning horizon. That is, at every period n , we determine a control policy for periods $\{n, n + 1, \dots, N\}$, but only implement the replenishment decision regarding period n .

It is important to remark that policy parameters of the dynamic uncertainty strategy are independent of the inventory position; i.e. the s and S of a non-stationary (s, S) policy remain optimal regardless of actual demand realisations. In this policy, demand realisations are accounted for at the beginning of each period, when inventory position is checked and, if it is found to be below s , it is raised to S by issuing an order of appropriate size. However, the same does not hold for static uncertainty and static-dynamic uncertainty strategies, for which the inventory position at the beginning of the planning horizon may affect the timing of replenishments that are scheduled in future periods, i.e. the R parameter in the (R, S) and (R, Q) policies. This observation motivates the need for re-planning in the context of (R, S) and (R, Q) policies, and hence the numerical study carried out in this section. Our aim is essentially to explore how the cost performance of different strategies is affected by the availability of realized demand information.

The receding horizon control applied in this study features similarities and differences from the conventional rolling horizon planning ([1], p. 199). In rolling horizon planning a fixed length of time window is rolled forward after the imminent period’s decision is implemented and the demand realization occurs. Rolling horizon planning for probabilistic demand is first implemented by Wemerlöv and Whybark [40] and then revisited by Bookbinder and H’ng [23] and Bookbinder and Tan [12]. These studies assume a long finite planning horizon so that fixed length of time window is rolled many times. As it is the case in the receding horizon control, towards the end of the horizon, these studies limit the length of time window with the number of remaining periods to capture end-of-horizon effects that are important for products featuring short life cycles. However, as the demand unfolds, these studies revise the demand forecasts for the rest of the planning horizon, in addition to updating the current inventory position.

All the approaches contrasted in our study operate under the assumption that demands in different periods are independently distributed; under this assumption, demand forecast update is clearly not justified. For this reason, in our study we simply update the current inventory position for re-planning purposes, leaving the demand forecasts unchanged.

Since, as we remarked, policy parameters of the dynamic uncertainty strategy are independent of the inventory level on hand,¹ in our study we only experiment on the receding horizon implemen-

¹ Note that this is not true if one updates demand forecasts at a given period, in which case new (s, S) levels will have to be recomputed for future periods since the problem instance has changed.

Fig. 1. Demand patterns.

tations of static uncertainty and static-dynamic uncertainty strategies.

In the following, we first present the experimental design adopted in our work, and then we discuss the results of the numerical study.

4.1. Experimental setup

Some of the studies under consideration do not immediately extend to generic demand distributions, e.g. the cumulative demand distribution of all possible replenishment cycles is required in [10,14,15], this convolution often cannot be obtained in closed form; similarly, Bollapragada and Morton [13] only discuss the case of a Poisson demand and of a normal demand with a fixed coefficient of variation across periods; as the authors remark, their approach can be conceptually generalised to other distributions, however this requires the definition of tailor made search strategies ([13], p. 578) and it is therefore not trivial. To enable a comparison across all approaches surveyed, demand d_n in each period n is therefore assumed to be an independently normally distributed random variable with known expected value $\tilde{d}_n$ and standard de-

viation $\sigma_n = \rho \cdot \tilde{d}_n$, where ρ denotes the coefficient of variation of the demand, which remains fixed over time as prescribed in [13]. As remarked, demand distributions are never updated, as this sort of action would be in contrast with the key assumption of demand independence over time periods, which is common across all methods we are working on.

We consider a planning horizon comprising 24 periods and 6 different patterns for the expected value of the demand in each period of the planning horizon: a stationary demand pattern (STAT); an erratic pattern (RAND); two life cycle patterns, one with lower (LCY1) and one with higher (LCY2) variation of the expected demand over the planning horizon; and finally two sinusoidal patterns, one with weaker (SIN1) and one with stronger (SIN2) oscillations. We adapted these demand patterns from Berry [41]; the same patterns have been extensively used in the literature [6,16,20,42]. Fig. 1 graphically illustrates all demand patterns analyzed in the numerical study. Note that the average demand per period is equal to 100 for all demand patterns. By setting the holding cost $h = 1$, and by varying the coefficient of variation $\rho \in \{0.10, 0.20, 0.30\}$, the fixed ordering cost $K \in \{250, 500, 1000, 2000\}$, and the penalty cost $b \in \{2, 5, 10\}$, we car-

Table 2
Average % optimality gaps of methods for different pivoting parameters.

	Dynamic		Static-Dynamic				Static	
	Ask	Bol	Tar	Tar-R	Ros	Ros-R	Sox/Var	Sox/Var-R
Demand pattern (π)								
STA	1.8	0.4	1.3	0.2	1.3	0.2	10.5	0.4
RAND	3.5	12.2	2.0	0.3	2.0	0.3	15.0	0.6
SIN1	2.8	1.9	1.5	0.2	1.4	0.2	10.7	0.4
SIN2	4.3	10.0	2.2	0.5	2.2	0.4	13.5	0.6
LCY1	4.6	1.9	1.2	0.1	1.1	0.1	12.8	0.4
LCY2	6.5	3.0	1.3	0.2	1.3	0.2	14.9	0.5
Coefficient of variation (ρ)								
0.1	3.6	6.0	0.2	0.1	0.2	0.0	5.2	0.1
0.2	4.0	4.8	1.2	0.2	1.2	0.2	12.4	0.4
0.3	4.2	3.9	3.3	0.5	3.3	0.4	21.2	0.9
Setup cost (K)								
250	3.3	4.3	2.0	0.4	2.0	0.4	25.0	0.7
500	3.1	4.3	1.8	0.3	1.8	0.3	14.3	0.5
1000	4.7	4.4	1.5	0.2	1.5	0.2	8.0	0.4
2000	4.5	6.6	0.9	0.1	0.9	0.1	4.3	0.3
Penalty cost (p)								
2	4.4	4.3	0.8	0.2	0.8	0.1	8.4	0.3
5	3.9	5.4	1.5	0.3	1.5	0.2	13.1	0.4
10	3.5	5.0	2.3	0.3	2.3	0.3	17.2	0.7
All instances								
AVG	3.9	4.9	1.6	0.3	1.6	0.2	12.9	0.5

ried out a full factorial analysis comprising a total of 216 test instances.

We first obtain the optimal parameters of dynamic uncertainty strategy for each test instance by means of the stochastic dynamic programming approach based on the functional equation in [7]. An open source implementation of the stochastic dynamic programming approach we used is available in Java² and Python.³ We adopted a discretization step of size one for the normal demand and a continuity correction factor of 0.5, this led to computational times in the order of 10^2 s for the stochastic dynamic programming approach. However, the reader should note that, since stochastic dynamic programming is pseudo-polynomial, an increase of the average value of the demand or of its standard deviation will lead to a dramatic increase of the state space boundaries and hence of computational times.

We solved all instances by using each of the methods presented in Table 1. We note that the heuristic proposed by Bollapragada and Morton [13] implicitly assumes an infinite planning horizon, and thus, it cannot account for the so-called end-of-horizon effect. In order to overcome this issue, the optimal policy parameters of dynamic uncertainty strategy are used for the last 8 periods of the planning horizon as it is done in their original study.

For methods adopting static uncertainty and static-dynamic uncertainty strategy (i.e. [15–17]), we also investigate the effect of a receding horizon control implementation.

We simulate the control policies obtained by implementing each of the aforementioned methods. We use the common random numbers simulation strategy (see e.g. [43]) and implement a stopping rule so as to achieve an estimation error of $\pm 0.1\%$ of the expected total cost with 0.95 confidence probability. Average costs are then compared against the average cost of the optimal policy [7] obtained by the stochastic dynamic program. The differences between the two are recorded as the percentage optimality gaps.

We carry out sensitivity analysis by fixing one parameter at a time—this is called the “pivot” parameter—while varying others over their respective domains. The results of the numerical study are summarized in Table 2 where the average optimality gap is re-

(a) All implementations

(b) Re-planning implementations – A closer look

Fig. 2. Boxplot comparison of methods considered.

ported for all problem instances sharing the same pivot parameter. Also, in Fig. 2(a) and in Fig. 2(b), we provide a comprehensive set of boxplots for conventional and receding horizon implementations of different methods. Note that the results regarding each method is given under the abbreviated name of the first author of the corresponding paper, i.e. Ask [10], Bol [13], Tar [16], Ros [17], and Sox/Var [14,15]; and the suffix “-R” is added to those where a receding horizon approach is adopted, i.e. Sox/Var-R, Tar-R, and

² <http://gwr3n.github.io/jsdp/>.

³ <https://pypi.python.org/pypi/inventoryanalytics/>.

Ros–R. Note that, Vargas [15] presents a special case implementation of Sox [14] by assuming a normally distributed demand. So, in the numerical setup of the current study these two are treated as a single method and abbreviated as Sox/Var.

4.2. Discussion

Below we first give a brief overview of the findings regarding each of the methods considered with and without receding horizon control. In particular, we draw conclusions on the effects of varying demand and cost parameters. We refer to the results reported in Table 2, Fig. 2(a) and (b).

The results indicate that Sox/Var – in the absence of a receding horizon control – provides the largest optimality gap. This result is rather consistent and it holds for all parameter settings yielding an average optimality gap of 13%. In Fig. 2(a) we can see that the optimality gap may reach 60% in the worst case. These results can mainly be attributed to the structure of the policy. Because both the timing and the size of replenishments are fixed at the beginning of the planning horizon, the static uncertainty strategy has no means to take recourse actions against demand uncertainty. Hence, uncertainty accumulates throughout the planning horizon and leads to a large optimality gap.

It is possible to observe that Tar and Ros, which adopt the static-dynamic uncertainty strategy, consistently yield the best cost performances among all methods considered, with an optimality gap around 2%. Furthermore, the deviation in the cost performance of these methods is rather low. This immediately shows that the static-dynamic uncertainty is very competitive despite the fact that it schedules replenishments in advance.

Ask and Bol yield optimality gaps around 4% and 5% respectively. These gaps are better than those produced by static uncertainty heuristics; however, they are larger than those produced by static-dynamic uncertainty heuristics. This finding is a particularly interesting one. The dynamic uncertainty strategy is by definition a superior strategy than the static-dynamic uncertainty strategy because it uses more information at the time of making replenishment decisions. However, this is true provided that both policies are implemented by using their respective optimal parameters. From our study it appears that existing heuristics for the dynamic uncertainty policy do not produce parameters of sufficient quality to outpace the performance of state-of-the-art static-dynamic uncertainty heuristics.

We now concentrate on the effects of varying demand and cost parameters. The results suggest, in general, that all methods perform better when the demand pattern is steady. This is particularly visible when we look at variants of the same underlying pattern characterized by small and large oscillations. The average optimality gap of all methods significantly increases as we move from SIN1 and SIN2, or LCY1 to LCY2. This immediately shows that managing inventories is more difficult when demand is heavily non-stationary.

We observe that performance of all methods deteriorates as demand variability increases. For instance, the optimality gap of Sox/Var gradually increases from 5.15% to 21.17% as the coefficient of variation increases from 0.1 to 0.3. The only exception of this is the heuristic of Bol. This method performs better under higher levels of demand uncertainty. This result can be attributed to the fact that all methods except Bol are adaptations of solution methods which were originally designed for the deterministic version of the lot sizing problem.

If we consider performance of all methods with respect to fixed ordering costs, we see that performance of “static” Sox/Var, as well as that of “static-dynamic” Tar and Ros improves when fixed ordering cost increases. This is not the case for “dynamic” uncertainty heuristics, i.e. Ask and Bol; the performance of these latter

heuristics does not seem to be affected by variation of the fixed ordering cost. This is probably due to the fact that a higher fixed ordering cost is likely to induce longer replenishment cycles; in turn, longer replenishment cycles mean fewer and more “stable” replenishments, in the sense that timing of a replenishment will hardly be affected by previous demand realisations. In this context, heuristics that seek a “static-dynamic” as opposed to a purely “dynamic” control are likely to produce a better outcome. When we look at the effect of stock-out penalty cost, both Sox/Var as well as the heuristics of Tar and Ros perform worse for higher values of penalty cost. Ask, on the other hand, performs better when penalty cost increases. Finally, the performance of the heuristic of Bol does not display a clear response to penalty cost variations, but remains worse than or comparable to Ask across all values considered.

Having summarized the results obtained from conventional implementations of different methods, we turn our attention on their receding horizon implementations. As remarked, we only target static-uncertainty and static-dynamic uncertainty strategies, since a receding horizon approach is only of value if the policy parameters are dependent on the inventory level on hand. Hence, we experiment on Sox/Var, Tar, and Ros. We observe that the receding horizon approach not only significantly improves the average cost performance of all these methods across all combinations of problem parameters considered, but it also reduces the dispersion in the worst-case performance (see Fig. 2(b)). In particular, we see that the heuristics of Tar and Ros perform extremely well when deployed under a receding horizon approach and both yield optimality gaps around 0.25%. On the other hand, probably the most remarkable result in our numerical experiments is that the receding horizon implementation of the method of Sox/Var adopting static uncertainty strategy displays an excellent performance with an optimality gap around 0.5% – slightly higher than those of the heuristics based on static-dynamic uncertainty strategy. Note that this method displays an average optimality gap of 13% without the receding horizon deployment. This clearly shows that the receding horizon approach makes it possible for the static uncertainty strategy to take suitable recourse actions against demand uncertainty.

4.3. Practical implications

We hereby reflect on practical implications of findings presented in our numerical study.

As pointed out in a number of works, see e.g. [5], to cope with demand uncertainty companies must use replenishment systems that adopt advanced inventory control policies. In the case of non-stationary demand—which is almost always the norm in industries facing seasonal patterns, trends, business cycles, and limited-life items—the structure of the optimal policy has been characterised in [7], and named “dynamic uncertainty” by Bookbinder and Tan [12]. However, adoption of this policy entails both implementation challenges, as well as computational challenges associated with the determination of optimal or near-optimal policy parameters. Practitioners may therefore benefit from the adoption of heuristic policies featuring a simpler structure.

Because of the complexity associated with the computation of optimal policy parameters, over the past decades research has mainly focused on the determination of near-optimal policy parameters, see e.g. [10,13]. Unfortunately, our numerical experiments demonstrate these heuristics may lead to large optimality gaps.

In the past, researchers have not explored the option of adopting heuristic policies, such as Bookbinder and Tan’s “static-uncertainty” or “static-dynamic uncertainty” [12] as an alternative to an optimal policy; these policies are very cheap to compute, and easy to implement, but unfortunately do not appear to be competitive with dynamic uncertainty heuristics Ask [10] and Bol [13].

Our findings, however, suggest that the picture changes under receding horizon control. In this setting, both static-uncertainty and static-dynamic uncertainty feature very competitive optimality gaps and fully dominate [10,13]. In turn, this means that focusing on approximating a control policy that follows the optimal form may not be advantageous in a heuristic setting, and that alternative policy structures should also be considered. Alternative policies may feature a simpler structure and may result easier to implement in practice. According to our study, those companies that do not have the required infrastructure to deal with the inherent complexity of optimal inventory management policies may adopt a static-uncertainty policy under a receding horizon control; this solution is cheaper to implement than other well-known heuristic policies (e.g. static-dynamic uncertainty), and features comparable optimality gaps. Finally, this form of control not only produces plans that feature competitive average optimality gaps, but that are also very reliable (i.e., low variance in the observed optimality gaps) and robust (i.e., stable performance over a wide-range of cost parameters).

These sought after features of the static-uncertainty policy under a receding horizon control makes it a promising inventory management solution for companies that (i) are exposed to non-stationary demand, (ii) cannot allocate extensive computational and human resources to inventory management, (iii) do not have the means to acquire sophisticated software for optimal control.

5. Conclusions

In this paper, we presented a cost-based performance evaluation of three lot-sizing strategies, i.e. static uncertainty, dynamic uncertainty, static-dynamic uncertainty, for non-stationary stochastic demand with backorder cost within a receding horizon framework. We first compared six heuristics: Ask [10], Bol [13], Tar [16], Ros [17], and Sox/Var [14,15] on a common test bed and computed their optimality gap with respect to the optimal control policy obtained via stochastic dynamic programming [7]. We, then carried out a similar cost comparison for a receding horizon deployment of the static and static-dynamic heuristics. We investigated the impact of the receding horizon scheme on the cost performances of separate inventory control policies.

According to our numerical study the static-dynamic uncertainty strategy is the one that displays the closest cost performance to the optimal dynamic uncertainty strategy. In particular, heuristic methods of Tar and Ros yield an average optimality gap of 2% as opposed to that of a static uncertainty strategy, which averages 13%. Surprisingly, heuristic methods proposed for dynamic uncertainty strategy yield a cost performance in-between the two other strategies. Furthermore, their cost performance is very sensitive to the parameter settings and displays large dispersion against the parameter changes. These results suggest that heuristics based on the static-dynamic uncertainty strategy are strong alternatives to the optimal lot-sizing policy.

When it comes to receding horizon deployment, however, the scene completely changes. As one would expect, receding horizon deployment improves the cost performance of all methods considered. Though, an outstanding impact is observed for the static uncertainty strategy, which now feature an average optimality gap of 0.5%. This shows that under receding horizon control the static uncertainty policy turns out to be one of the strongest alternatives to the optimal dynamic uncertainty strategy.

This study emphasises the fact that receding horizon control leverages realised demand information and gives the opportunity to take appropriate recourse actions which enhance performance of policies, such as the static uncertainty policy, that would otherwise perform poorly. These recourse actions, however, improve cost performance at the expense of higher sys-

tem nervousness. Provided a decision maker is willing to accept this increased system nervousness, a well designed receding horizon implementation provides a competitive cost performances in non-stationary stochastic inventory control under a finite planning horizon.

References

- [1] Silver EA, Pyke DF, Peterson R. Inventory management and production planning and scheduling. New York: Wiley; 1998.
- [2] Cui AS, Zhao M, Ravichandran T. Market uncertainty and dynamic new product launch strategies: a system dynamics model. *IEEE Trans Eng Manag* 2011;58(3):530–50.
- [3] Kaki A, Salo A, Talluri S. Scenario-based modeling of interdependent demand and supply uncertainties. *IEEE Trans Eng Manag* 2014;61(1):101–13.
- [4] Lu H, Wang H, Xie Y, Li H. Construction material safety-stock determination under nonstationary stochastic demand and random supply yield. *IEEE Trans Eng Manag* 2016;63(2):201–12.
- [5] Graves SC, Willems SP. Strategic inventory placement in supply chains: non-stationary demand. *Manuf Serv Oper Manag* 2008;10(2):278–87.
- [6] Tunc H, Kilic OA, Tarim SA, Eksioglu B. The cost of using stationary inventory policies when demand is non-stationary. *Omega* 2011;39(4):410–15.
- [7] Scarf HE. Optimality of (s, S) policies in the dynamic inventory problem. In: Arrow K, Karlin S, Suppes P, editors. *Mathematical methods in the social sciences*. Stanford, CA: Stanford University Press; 1960. p. 196–202.
- [8] Veinott AF, Wagner HM. Computing optimal (s, S) inventory policies. *Manag Sci* 1965;11(5):525–52.
- [9] Tunc H, Kilic OA, Tarim SA, Eksioglu B. A simple approach for assessing the cost of system nervousness. *Int J Prod Econ* 2013;141(2):619–25.
- [10] Askin RG. A procedure for production lot sizing with probabilistic dynamic demand. *IIE Trans* 1981;13(2):132–7.
- [11] Silver EA, Meal HA. A heuristic for selecting lot size requirements for the case of a deterministic time varying demand rate with discrete opportunities for replenishment. *Prod Inventory Manag J* 1973;14(2):63–73.
- [12] Bookbinder JH, Tan JY. Strategies for the probabilistic lot-sizing problem with service-level constraints. *Manag Sci* 1988;34(9):1096–108.
- [13] Bollapragada S, Morton T. A simple heuristic for computing non-stationary (s, S) policies. *Oper Res* 1999;47(4):576–84.
- [14] Sox CR. Dynamic lot sizing with random demand and non-stationary costs. *Oper Res Lett* 1997;20(4):155–64.
- [15] Vargas V. An optimal solution for the stochastic version of the Wagner–Whitin dynamic lot-size model. *Eur J Oper Res* 2009;198(2):447–51.
- [16] Tarim SA, Kingsman BG. Modelling and computing (R^n, S^n) policies for inventory systems with non-stationary stochastic demand. *Eur J Oper Res* 2006;174(1):581–99.
- [17] Rossi R, Tarim SA, Hnich B, Prestwich S. Constraint programming for stochastic inventory systems under shortage cost. *Ann Oper Res* 2012;195(1):49–71.
- [18] Tunc H, Kilic OA, Tarim SA, Rossi R. An extended mixed-integer programming formulation and dynamic cut generation approach for the stochastic lot-sizing problem. *INFORMS J Comput* 2018;30(3):492–506.
- [19] Tarim SA, Kingsman BG. The stochastic dynamic production/inventory lot-sizing problem with service-level constraints. *Int J Prod Econ* 2004;88(1):105–19.
- [20] Rossi R, Tarim SA, Hnich B, Prestwich S. A global chance-constraint for stochastic inventory systems under service level constraints. *Constraints* 2008;13(4):490–517.
- [21] Tarim SA, Dogru MK, Ozen U, Rossi R. An efficient computational method for a stochastic dynamic lot-sizing problem under service-level constraints. *Eur J Oper Res* 2011;215(3):563–71.
- [22] Rossi R, Kilic OA, Tarim SA. Piecewise linear approximations for the static-dynamic uncertainty strategy in stochastic lot-sizing. *Omega* 2015;50:126–40.
- [23] Bookbinder JH, H'ng BT. Production lot sizing for deterministic rolling schedules. *J Oper Manag* 1986;6(3):349–62.
- [24] Kwon WH, Han SH. Receding horizon control: model predictive control for state models. Springer Science & Business Media; 2006.
- [25] Kadipasaoglu S, Sridharan S. Measurement of instability in multi-level mrp systems. *Int J Prod Res* 1997;35(3):713–37.
- [26] Kilic OA, Tarim SA. An investigation of setup instability in non-stationary stochastic inventory systems. *Int J Prod Econ* 2011;133(1):286–92.
- [27] Wang Z, Chan FT. A robust replenishment and production control policy for a single-stage production/inventory system with inventory inaccuracy. *IEEE Trans Syst Man Cybern* 2015;45(2):326–37.
- [28] Wang Z, Chan FT, Li M. A robust production control policy for hedging against inventory inaccuracy in a multiple-stage production system with time delay. *IEEE Trans Eng Manag* 2018;65(3):474–86. doi:10.1109/tem.2018.2792319.
- [29] Wagner HM, Whitin TM. Dynamic version of the economic lot size model. *Manag Sci* 1958;5(12):89–96.
- [30] Ehrhardt R. Power approximation for computing (s, S) inventory policies. *Manag Sci* 1979;25(8):777–86.
- [31] Federgruen A, Zipkin P. An efficient algorithm for computing optimal (s, S) policies. *Oper Res* 1984;32(6):1268–85.
- [32] Zheng YS, Federgruen A. Finding optimal (s, S) policies is about as simple as evaluating a single policy. *Oper Res* 1991;39(4):654–65.
- [33] Kingsman BG. Raw materials purchasing: an operational research approach, 4. 1st ed. Pergamon Press; 1985.

- [34] Li Q, Wu X, Cheung KL. Optimal policies for inventory systems with separate delivery-request and order-quantity decisions. *Oper Res* 2009;57(3):626–36.
- [35] Mutlu F, Cetinkaya S, Bookbinder JH. An analytical model for computing the optimal time-and-quantity-based policy for consolidated shipments. *IIE Trans* 2010;42:367–77.
- [36] Silver EA. Inventory control under a probabilistic time-varying, demand pattern. *IIE Trans* 1978;10(4):371–9.
- [37] Focacci F, Lodi A, Milano M. Optimization-oriented global constraints. *Constraints* 2002;7(3):351–65.
- [38] Rossi R, Tarim S, Hnich B, Prestwich S. Cost-based domain filtering for stochastic constraint programming. In: *Principles and practice of constraint programming*. Springer; 2008. p. 235–50.
- [39] Rossi R, Kilic OA, Tarim SA. Piecewise linear approximations for the static-dynamic uncertainty strategy in stochastic lot-sizing. *Omega* 2015;50:126–140.
- [40] Wemmerlöv U, Whybark DC. Lot-sizing under uncertainty in a rolling schedule environment. *Int J Prod Res* 1984;22(3):467–84.
- [41] Berry WL. Lot sizing procedures for requirements planning systems: a framework for analysis. *Prod Inventory Manag J* 1972;13(2):19–34.
- [42] Ho JC, Solis AO, Chang YL. An evaluation of lot-sizing heuristics for deteriorating inventory in material requirements planning systems. *Comput Oper Res* 2007;34(9):2562–75.
- [43] Sloan BJ, Unwin AR. Common random numbers in multivariate simulations. *Eur J Oper Res* 1990;48(2):252–9.