

Title	The benefit of receding horizon control: near-optimal policies for stochastic inventory control
Authors	Dural-Selcuk, Gozdem;Rossi, Roberto;Kilic, Onur A.;Tarim, S. Armagan
Publication date	2019-07-24
Original Citation	Dural-Selcuk, G., Rossi, R., Kilic, O. A. and Tarim, S. A. (2019) 'The benefit of receding horizon control: near-optimal policies for stochastic inventory control', Omega. doi: 10.1016/j.omega.2019.07.007
Type of publication	Article (peer-reviewed)
Link to publisher's version	http://www.sciencedirect.com/science/article/pii/S0305048319302440 - 10.1016/j.omega.2019.07.007
Rights	© 2019, Elsevier Ltd. All rights reserved. This manuscript version is made available under the CC BY-NC-ND 4.0 licence. - https://creativecommons.org/licenses/by-nc-nd/4.0/
Download date	2024-04-24 14:14:54
Item downloaded from	https://hdl.handle.net/10468/8409

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

^a*Industrial Engineering, Atılım University, Ankara, Turkey*

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

^c*Faculty of Economics and Business, University of Groningen, Groningen, The Netherlands*

^d*Cork University Business School, University College Cork, Ireland*

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.

Keywords: stochastic lot sizing, static uncertainty, dynamic uncertainty, static-dynamic uncertainty, receding horizon control

*Corresponding author

Email addresses: gozdemduralselfcuk@gmail.com (Gozdem Dural-Selcuk), roberto.rossi@ed.ac.uk (Roberto Rossi), o.a.kilic@rug.nl (Onur A. Kilic), armagan.tarim@ucc.ie (S. Armagan Tarim)

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] and [13]. The framework introduced by [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, 15] and [16, 17]. 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 [18, 19, 20] and β (i.e. fill rate) service level constraints [21]. 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 [22], Bookbinder and H'ng compared the performance of six separate lot-sizing methods for use in a rolling schedule. To the best of our knowledge, there is no similar 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 strategies over different solution approaches for the single-item, single-stocking point, non-stationary stochastic lot-sizing problem under backorder costs;
- We, for the first time, compare the cost performance of these solution approaches in a receding horizon deployment [23], which is widely used in the lot-sizing literature [i.e. 24, 25]. 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 some managerial insights; displaying 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

60 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
 65 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
 70 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]:

$$\min \sum_{n=1}^N \mathbb{E} \{ K z_n + h x_n^+ + b x_n^- \} \quad (1)$$

75 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
 80 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. [26], [27]).

2.1. Control policies

85 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 90 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 95 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 .

100 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 105 inventory level.

[9] numerically analyse the cost performance of these strategies. They 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 110 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 [28]. 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

140 be considered. A variety of studies address this issue [see e.g. 8, 29, 30, 31].
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. [10] proposes a heuristic method which
145 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
150 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
155 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.
[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
160 follows. First, by means of the method developed by [31], 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
165 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
170 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 replenish-
175 ments at the beginning of the planning horizon. It is customary in the liter-
ature 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
180 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 ob-
served. Despite being unable to exploit additional information obtained from
past demand realisations, the static uncertainty strategy has the advantage of
185 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 formula-
tion of the problem where each arc corresponds to a prospective replenishment
190 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 [28] with additional
195 feasibility constraints. [15] employs the same approach 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 nor-
mally distributed, and provides a simple, yet very efficient solution algorithm
by exploiting the properties of the normal distribution.

200 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
205 “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 [25, 9]. Because
210 it offers a fixed order schedule, the static-dynamic strategy is particularly appealing in material requirement planning, joint replenishment, and shipment consolidation environments [see 32, 1, 33, 34].

The contributions in this line of research are as follows. [35] proposes a heuristic which can be regarded as the stochastic version of [11]. As is the case in
215 [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
220 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
225 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 [35], their method determines order-periods and order-up-to levels simultaneously under
230 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
 235 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,
 240 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. [17] develop a stochastic constraint programming model which is mainly equivalent to the MIP model in [16]. However, instead of employing
 245 a piecewise linear approximation, they use the exact non-linear cost function. They also introduce a cost based filtering method [see e.g. 36, 37] 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
 250 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
 255 inventories in cases where the actual inventory level exceeds the order-up-to level. More recently, [38] 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
 260

Table 1: An overview of the methods considered in the current study; “Comb.” denotes a combinatorial complexity.

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	Comb. (MIP)
Rossi et al. [17]	Ros	Static-Dynamic	Heuristic	Comb. (MIP)
Sox [14]	Sox	Static	Optimal	$O(N^2)$ (DP)
Vargas [15]	Var	Static	Optimal	Comb. (MIP)

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 . [16], [17] and [15] 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
285 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 [23], which is widely used in the lot-sizing literature [i.e. 24, 25]. The receding horizon control in inventory problems proceeds as follows: the replenishment plans are made over the entire planning
290 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
295 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 ,
300 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
305 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
310 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 [39] and then revisited by [22] and [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 implementations 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*

We consider a planning horizon comprising 24 periods. Some of the studies under consideration do not immediately extend to generic demand distributions, e.g. [10, 14, 15] requires the cumulative demand distribution of all possible

¹Note 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.

replenishment cycles, a convolution that often cannot be obtained in closed form; similarly, [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
340 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 $n = 1, \dots, 24$ is therefore assumed to be an independently normally distributed random variable with known expected value \tilde{d}_n and standard deviation
345 $\sigma_n = \rho \cdot \tilde{d}_n$, 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
350 across all methods we are working on.

We consider 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
355 finally two sinusoidal patterns, one with weaker (SIN1) and one with stronger (SIN2) oscillations. We adapted these demand patterns from [40]; the same patterns have been extensively used in the literature [16, 41, 19, 6]. Figure 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.
360 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 generate 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
365 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 seconds for the stochastic dynamic programming approach.

370 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.

375 We note that the heuristic proposed by [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.

380 For methods adopting static uncertainty and static-dynamic uncertainty strategy [i.e. 15, 16, 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

385 [see e.g. 42] 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.

390 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 reported for all problem instances characterized by

²<http://gwr3n.github.io/jsdp/>

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

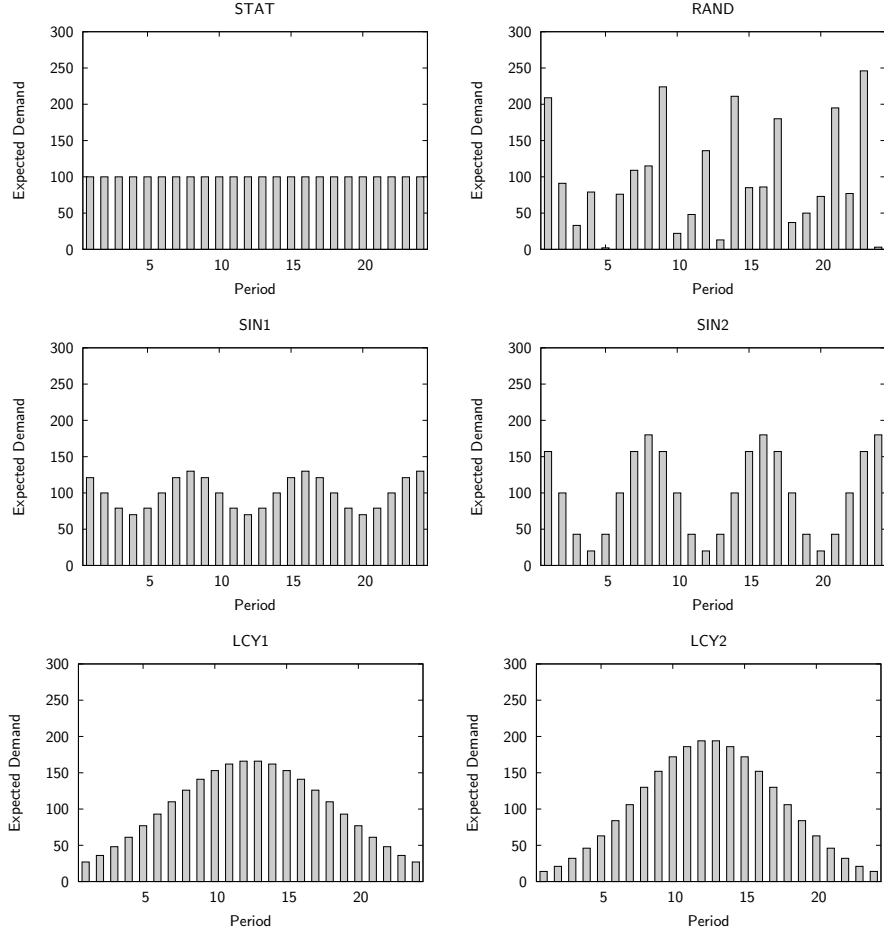


Figure 1: Demand patterns

the same pivot parameter. Also, in Figure 2(a) and in Figure 2(b), we provide
 395 a comprehensive set of boxplots for conventional and receding horizon imple-
 mentations 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.
 400 Sox/Var-R, Tar-R, and 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

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


The results indicate that Sox/Var – in the absence of a receding horizon
410 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 Figure 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
415 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
420 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.


425 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
430 because it uses more information at the time of making replenishment decisions. However, this is true provided optimal policy parameters are available for both

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



(a) All implementations



(b) Re-planning implementations – A closer look

Figure 2: Boxplot comparison of methods considered

policies. 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.

435 We now concentrate on the effects of varying demand and cost parameters. The results suggest, in general, that all methods perform relatively better when the demand pattern is rather steady. This is especially apparent when we look at variants of the same underlying pattern characterized by small and large oscillations. The average optimality gap of all methods significantly increase as
 440 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 relatively better in higher levels of demand uncertainty. This result can mainly 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 likely 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 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 relatively better as the 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 Figure 2(b)). In particular, we see that the heuristics of Tar and Ros perform extremely well when deployed with 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 sophisticated inventory control policies. In the case of non-stationary demand — which is almost always the norm in industries where seasonal patterns, trends, business cycles, and limited-life items are observed — the structure of the optimal policy has been characterised in [7], and named “dynamic uncertainty” by [12]. However, adoption of this policy is prohibitive in practice due to its implementation complexity and associated computational challenges. This in turn implies that practitioners must resort to heuristic policies, since seeking optimality would rule out any chance of successful implementation.

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
505 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 uncer-
510 tainty 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.
515 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 (SU-RH) control; this solution is cheaper to implement than other well-known heuristic policies (e.g. static-dynamic uncertainty), and features comparable optimality
520 gaps. Finally, SU-RH 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 SU-RH makes it a promising inventory
525 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

530 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 firstly 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 optimality gap around 2% as opposed to the large optimality gap of static uncertainty strategy around 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 the heuristics based on static-dynamic uncertainty strategy are strong alternatives for the optimal lot-sizing policy in the conventional implementation.

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

This study emphasizes the fact that receding horizon control incorporates realized demand information and gives the opportunity to take recourse actions for the poorly performing policies when applied in conventional form. The recourse actions, however, improves the cost performance against the backdrop of higher system nervousness. Despite this, a well designed receding horizon

implementation would still provide a fair trade off with the advantage of better
 cost performances in non-stationary stochastic inventory problems under a finite
 565 planning horizon.

References

- [1] E. A. Silver, D. F. Pyke, R. Peterson, Inventory Management and Production Planning and Scheduling, Wiley New York, 1998.
- [2] A. S. Cui, M. Zhao, T. Ravichandran, Market uncertainty and dynamic new
 570 product launch strategies: a system dynamics model, IEEE Transactions
 on Engineering Management 58 (3) (2011) 530–550.
- [3] A. Kaki, A. Salo, S. Talluri, Scenario-based modeling of interdependent
 demand and supply uncertainties, IEEE Transactions on Engineering Man-
 agement 61 (1) (2014) 101–113.
- [4] H. Lu, H. Wang, Y. Xie, H. Li, Construction material safety-stock determi-
 575 nation under nonstationary stochastic demand and random supply yield,
 IEEE Transactions on Engineering Management 63 (2) (2016) 201–212.
- [5] S. C. Graves, S. P. Willems, Strategic inventory placement in supply chains:
 Nonstationary demand, Manufacturing & Service Operations Management
 10 (2) (2008) 278–287.
 580
- [6] H. Tunc, O. A. Kilic, S. A. Tarim, B. Eksioglu, The cost of using stationary
 inventory policies when demand is non-stationary, Omega 39 (4) (2011)
 410–415.
- [7] H. E. Scarf, Optimality of (s, S) policies in the dynamic inventory problem,
 585 in: K. Arrow, S. Karlin, P. Suppes (Eds.), Mathematical Methods in the
 Social Sciences, Stanford University Press, Stanford, CA, 1960, pp. 196–
 202.
- [8] A. F. Veinott, H. M. Wagner, Computing optimal (s, S) inventory policies,
 Management Science 11 (5) (1965) 525–552.


- 590 [9] H. Tunc, O. A. Kilic, S. A. Tarim, B. Eksioglu, A simple approach for assessing the cost of system nervousness, *International Journal of Production Economics* 141 (2) (2013) 619–625.
- [10] R. G. Askin, A procedure for production lot sizing with probabilistic dynamic demand, *IIE Transactions* 13 (2) (1981) 132–137.
- 595 [11] E. A. Silver, H. A. Meal, A heuristic for selecting lot size requirements for the case of a deterministic time varying demand rate with discrete opportunities for replenishment, *Production and Inventory Management Journal* 14 (2) (1973) 63–73.
- [12] J. H. Bookbinder, J. Y. Tan, Strategies for the probabilistic lot-sizing problem with service-level constraints, *Management Science* 34 (9) (1988) 1096–
600 1108.
- [13] S. Bollapragada, T. Morton, A simple heuristic for computing non-stationary (s, S) policies, *Operations Research* 47 (4) (1999) 576–584.
- [14] C. R. Sox, Dynamic lot sizing with random demand and non-stationary costs, *Operations Research Letters* 20 (4) (1997) 155–164.
605
- [15] V. Vargas, An optimal solution for the stochastic version of the Wagner–Whitin dynamic lot-size model, *European Journal of Operational Research* 198 (2) (2009) 447–451.
- [16] S. A. Tarim, B. G. Kingsman, Modelling and computing (R^n, S^n) policies
610 for inventory systems with non-stationary stochastic demand, *European Journal of Operational Research* 174 (1) (2006) 581–599.
- [17] R. Rossi, S. A. Tarim, B. Hnich, S. Prestwich, Constraint programming for stochastic inventory systems under shortage cost, *Annals of Operations Research* 195 (1) (2012) 49–71.
- 615 [18] S. A. Tarim, B. G. Kingsman, The stochastic dynamic production/inventory lot-sizing problem with service-level constraints, *International Journal of Production Economics* 88 (1) (2004) 105 – 119.

- [19] R. Rossi, S. A. Tarim, B. Hnich, S. Prestwich, A global chance-constraint for stochastic inventory systems under service level constraints, *Constraints* 13 (4) (2008) 490–517.
- [20] S. A. Tarim, M. K. Dogru, U. Ozen, R. Rossi, An efficient computational method for a stochastic dynamic lot-sizing problem under service-level constraints, *European Journal of Operational Research* 215 (3) (2011) 563–571.
- [21] R. Rossi, O. A. Kilic, S. A. Tarim, Piecewise linear approximations for the static–dynamic uncertainty strategy in stochastic lot-sizing, *Omega* 50 (2015) 126–140.
- [22] J. H. Bookbinder, B.-T. H'ng, Production lot sizing for deterministic rolling schedules, *Journal of Operations Management* 6 (3) (1986) 349–362.
- [23] W. H. Kwon, S. H. Han, Receding horizon control: model predictive control for state models, Springer Science & Business Media, 2006.
- [24] S. Kadipasaoglu, S. Sridharan, Measurement of instability in multi-level mrp systems, *International Journal of Production Research* 35 (3) (1997) 713–737.
- [25] O. A. Kilic, S. A. Tarim, An investigation of setup instability in non-stationary stochastic inventory systems, *International Journal of Production Economics* 133 (1) (2011) 286–292.
- [26] Z. Wang, F. T. Chan, A robust replenishment and production control policy for a single-stage production/inventory system with inventory inaccuracy, *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 45 (2) (2015) 326–337.
- [27] Z. Wang, F. T. Chan, M. Li, A robust production control policy for hedging against inventory inaccuracy in a multiple-stage production system with time delay, *IEEE Transactions on Engineering Management*.


- [28] H. M. Wagner, T. M. Whitin, Dynamic version of the economic lot size
645 model, *Management Science* 5 (12) (1958) 89–96.
- [29] R. Ehrhardt, Power approximation for computing (s, S) inventory policies,
Management Science 25 (8) (1979) 777–786.
- [30] A. Federgruen, P. Zipkin, An efficient algorithm for computing optimal
 (s, S) policies, *Operations Research* 32 (6) (1984) 1268–1285.
- 650 [31] Y. S. Zheng, A. Federgruen, Finding optimal (s, S) policies is about as
simple as evaluating a single policy, *Operations Research* 39 (4) (1991)
654–665.
- [32] B. G. Kingsman, *Raw materials purchasing: An operational research ap-
proach*, 1st Edition, Vol. 4, Pergamon Press, 1985.
- 655 [33] Q. Li, X. Wu, K. L. Cheung, Optimal policies for inventory systems
with separate delivery-request and order-quantity decisions, *Operations Re-
search* 57 (3) (2009) 626–636.
- [34] F. Mutlu, S. Cetinkaya, J. H. Bookbinder, An analytical model for comput-
ing the optimal time-and-quantity-based policy for consolidated shipments,
660 *IIE Transactions* 42 (2010) 367–377.
- [35] E. A. Silver, Inventory control under a probabilistic time-varying, demand
pattern, *IIE Transactions* 10 (4) (1978) 371–379.
- [36] F. Focacci, A. Lodi, M. Milano, Optimization-oriented global constraints,
Constraints 7 (3) (2002) 351–365.
- 665 [37] R. Rossi, S. Tarim, B. Hnich, S. Prestwich, Cost-based domain filtering
for stochastic constraint programming, in: *Principles and Practice of Con-
straint Programming*, Springer, 2008, pp. 235–250.
- [38] R. Rossi, O. A. Kilic, S. A. Tarim, Piecewise linear approximations for
the static-dynamic uncertainty strategy in stochastic lot-sizing, *Omega* 50
670 (2015) 126–140.

- [39] U. Wemmerlöv, D. C. Whybark, Lot-sizing under uncertainty in a rolling schedule environment, *The International Journal Of Production Research* 22 (3) (1984) 467–484.
- [40] W. L. Berry, Lot sizing procedures for requirements planning systems: A
675 framework for analysis, *Production and Inventory Management Journal* 13 (2) (1972) 19–34.
- [41] J. C. Ho, A. O. Solis, Y.-L. Chang, An evaluation of lot-sizing heuristics for deteriorating inventory in material requirements planning systems, *Computers & Operations Research* 34 (9) (2007) 2562–2575.
- [42] B. J. Sloan, A. R. Unwin, Common random numbers in multivariate simu-
680 lations, *European Journal of Operational Research* 48 (2) (1990) 252–259.


Optimality Gap (%)




Optimality Gap (%)




LCY1




LCY2




RAND



SIN1



SIN2



STAT

