

Title	A mathematical programming-based solution method for the nonstationary inventory problem under correlated demand
Authors	Xiang, Mengyuan;Rossi, Roberto;Martin-Barragan, Belen;Tarim, S. Armagan
Publication date	2022-04-22
Original Citation	Xiang, M., Rossi, R., Martin-Barragan, B. and Tarim, S. A. (2022) 'A mathematical programming-based solution method for the nonstationary inventory problem under correlated demand', European Journal of Operational Research, 304(2), pp.515-524. https://doi.org/10.1016/j.ejor.2022.04.011
Type of publication	Article (peer-reviewed)
Link to publisher's version	10.1016/j.ejor.2022.04.011
Rights	© 2022, Elsevier B.V. All rights reserved. This manuscript version is made available under the CC BY-NC-ND 4.0 license. - https://creativecommons.org/licenses/by-nc-nd/4.0/
Download date	2026-09-20 06:46:17
Item downloaded from	https://hdl.handle.net/10468/13214

A mathematical programming-based solution method for the nonstationary inventory problem under correlated demand

Mengyuan Xiang^{a,*}, Roberto Rossi^b, Belen Martin-Barragan^b, S. Armagan
Tarim^c

^a*School of Intelligent Finance and Business, Entrepreneur College, Xi'an
Jiaotong-Liverpool University, Suzhou, China*

^b*Business School, University of Edinburgh, Edinburgh, United Kingdom*

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

Abstract

This paper extends the single-item single-stocking location nonstationary stochastic inventory problem to relax the assumption of independent demand. We present a mathematical programming-based solution method built upon an existing piecewise linear approximation strategy under the receding horizon control framework. Our method can be implemented by leveraging off-the-shelf mixed-integer linear programming solvers. It can tackle demand under various assumptions: the multivariate normal distribution, a collection of time-series processes, and the Martingale Model of Forecast Evolution. We compare against exact solutions obtained via stochastic dynamic programming to demonstrate that our method leads to near-optimal plans.

Keywords: inventory, correlated demand, stochastic programming, mixed integer linear programming, martingale model of forecast evolution

*Corresponding author. +86(0)512 8818 9024

Email addresses: Mengyuan.Xiang@xjtlu.edu.cn (Mengyuan Xiang), roberto.rossi@ed.ac.uk (Roberto Rossi), Belen.Martin@ed.ac.uk (Belen Martin-Barragan), armagan.tarim@ucc.ie (S. Armagan Tarim)

1. Introduction

Demand process characteristic plays an important role in determining the optimal inventory policy in an inventory model (Sethi & Cheng, 1997). Neale & Willems (2009) argued that a nonstationary stochastic demand is the rule, rather than the exception. Many early work focused on modelling nonstationary stochastic inventory problems under the assumption that each period's demand is independent to the demand in other periods. However, as discussed in Song & Zipkin (1993), environmental factors, such as economic conditions, market conditions, and any exogenous conditions, have significant effects on the demand for a product, the supply, and the cost structure. In this respect, this paper aims to *relax the assumption of demand independence between time periods*.

The literature on correlated demand can be roughly categorised into two groups. The first group assumes that the demand is driven by exogenous random factors, such as economic conditions and market conditions. The Markovian model is commonly adopted (Song & Zipkin, 1993; Sethi & Cheng, 1997; Beyer & Sethi, 1997), when the decision-makers fully observe demand states, and true demand distributions are known. The second group assumes demand distributions are updated with internal information, such as past demands; to capture this, time-series processes are adopted (Johnson & Thompson, 1975; Graves, 1999). This study falls into the second group, in which demand forecasts are updated based on past demands. Recently, the Martingale Model of Forecast Evolution (MMFE) has been extensively investigated, as it captures a wide range of forecasting techniques, such as time-series processes and judgement-based methods (Norouzi & Uzsoy, 2014; Albey et al., 2015). The model proposed in this study can be adapted to tackle demand updates following the MMFE.

We consider the periodic-review single-item single-stocking location nonstationary inventory problem under fixed and unit ordering costs, holding costs, and penalty costs. We assume that demands in different periods follow a multivariate normal distribution and that demands in successive periods are corre-

lated. After observing demand realisations, demand forecasts for the remaining periods are updated. Replenishment plans that use these forecasts take into account the evolution of these updates over time.

This study presents a novel approach for tackling the nonstationary lot-sizing problem under correlated demand. We leverage two key building blocks: modelling techniques originally discussed in Rossi et al. (2015), and results from multivariate stochastic process analysis. Rossi et al. (2015) presented a mixed-integer linear programming (MILP) model for approximating the optimal (R, S) policies under the assumption of independently distributed demand. The (R, S) policy fixes the timing of inventory reviews (R) and associated order-up-to-levels (S) at the beginning of the planning horizon; actual order quantities are decided upon only at the beginning of each inventory review period. We extend the existing modelling technique of Rossi et al. (2015), and we combine it with results from multivariate stochastic process analysis. The resulting model can tackle nonstationary stochastic inventory problems under demand following a multivariate normal distribution. The model developed is operationalized under a receding horizon control framework, in which only the imminent replenishment plan is implemented, and a re-planning is done at the beginning of each period for the rest of the planning horizon. Besides, we show that our approach can be adapted to tackle demand following a collection of time-series process and the MMFE.

Our work presents the first approach that combines exiting modelling techniques with results from multivariate stochastic process analysis for tackling nonstationary stochastic inventory problems under correlated demand in the context of a receding horizon control. It comes with the advantage of being able to tackle demands following the multivariate normal distribution, a collection of time-series processes, and the MMFE.

Our contributions to the literature on stochastic lot-sizing are the following.

- We present the first mathematical programming-based solution method for tackling the nonstationary stochastic inventory problem under correlated

demand.

- We combine an existing piecewise linear approximation strategy with results from multivariate normal probability theory in the context of a receding horizon control framework to model and solve the problem via a mixed integer linear programming model that can be solved by using off-the-shelf software.
- Our approach can tackle correlated demand under various assumptions: the multivariate normal distribution, a collection of time-series processes, and the MMFE.
- We compare our approach against solutions obtained via stochastic dynamic programming, and demonstrate that our approach leads to near-optimal solutions, while being able to tackle larger instances.

The rest of this paper is organised as follows. Section 2 surveys literature on stochastic inventory problems under demand following Markovian processes, time-series processes, and the MMFE. Section 3 discusses multivariate normal probability theories and the stochastic dynamic programming (SDP) formulation under correlated demand. Section 4 illustrates a stochastic optimization model of the (R, S) policy. Section 5 introduces our solution method for tackling the nonstationary stochastic lot-sizing problem under correlated demand. Section 6 presents our computational study. Finally, we draw conclusions in Section 7.

2. Literature review

In his landmark study, Scarf (1960) proved the optimality of (s, S) policy for the stochastic inventory problem with independent demand. Many researchers attempted to relax the assumption of independence by considering Markovian demand. Sethi & Cheng (1997) considered the discrete-time finite-horizon stochastic inventory problem with Markovian demand. They showed

that a state-dependent (s, S) policy is optimal in the context of total cost minimisation, which consists of fixed and proportional ordering costs, proportional holding costs, and proportional backorder costs. Under the same cost structure, Beyer & Sethi (1997) proved the optimality of (s, S) policy from the viewpoint of long-run average cost minimisation. Beyer et al. (1998) generalised discussions in (Sethi & Cheng, 1997; Beyer & Sethi, 1997), and proved that the (s, S) -type policy is optimal for the finite-horizon problem, discounted-cost infinite-horizon problem, and long-run average-cost problem. Cheng & Sethi (1999) further extended the Markovian demand models to incorporate the case in which the unsatisfied demand is lost rather than backlogged, and showed that the (s, S) policy is optimal. In contrast to existing studies that focus on proving the optimality of time-dependent (s, S) policy and computing the optimal expected costs, Nasr & Elshar (2018) presented a computational framework utilising a Markovian representation to evaluate the performance measures, which include the number of backorders, on-hand inventory, inventory position, and the ordering count process of the inventory system.

For the continuous-review infinite-horizon inventory problem with Markov-modulated demand, Song & Zipkin (1993) proved that the state-dependent base-stock policy is optimal when the order cost is linear in the order quantity, and the state-dependent (s, S) policy is optimal if there is a fixed ordering cost. An exact procedure and a modified value-iteration algorithm were presented to compute the optimal policy parameters. Lian et al. (2009) studied a continuous-review model for items with an exponential random lifetime and a general Markovian renewal demand process. They derived an analytical expression for computing the expected time between successive orders, the expected minimised long-run average cost, and the (s, S) policy parameters. Hu et al. (2016) investigated the periodic-review infinite-horizon inventory system with a Markov-modulated demand process under (s, S) policy. They proposed a Maclaurin series analysis and a Pade approximation for computing the optimal policy parameters.

The collective insight of these studies is that the optimal policy, either the (s, S) policy or the base-stock policy, for a model is state-dependent to reflect

the dynamics of the underlying demand environment. Demand states are fully observed by decision-makers, and true demand distributions are known. This study differs from existing works by assuming that future demand distributions depend on past demands, and thus demand states are unknown until past demands are observed.

Early research on demand forecast updates generally adopted time-series processes, which use internal information, such as past demands, to update future demand distributions. Johnson & Thompson (1975) proved the optimality of the base-stock policy for the single-item periodic-review inventory system with proportional holding and stock-out costs and zero lead time under a mixed autoregressive-moving average demand process and the condition that demands fall in a certain lower and upper bounds. Ray (1981) derived an analytical expression for computing the reorder level for which the demand follows the autoregressive (AR) and moving average (MA) processes. Under the same demand processes, Fotopoulos et al. (1988) presented a straightforward method based on the basis of probability arguments to approximate the reorder point and safety stock for which the lead time is arbitrary. Graves (1999) computed the base-stock policy for a single inventory system where the demand follows an integrated moving average (IMA) process. These existing studies tackle stochastic inventory problems under the assumption that demand follows certain types of time-series models. However, our method generalises existing studies and can tackle demand following a collection of time-series processes.

The Martingale Model of Forecast Evolution (MMFE) has recently been extensively explored in the modelling of correlated demand as it captures a wide range of forecasting techniques such as time-series processes and judgement-based methods. The MMFE assumes that demand forecasts for a number of periods in the future evolve over time as new information becomes available in each period. This framework was proposed by (Graves et al., 1986; Heath & Jackson, 1994) under the assumption that the information available to make forecasts grows as time grows, forecast updates are mean zero and uncorrelated with past observations, and forecast updates are stationary.

The MMFE has been extended and implemented by a number of authors in the context of inventory control models. Toktay & Wein (2001) analysed the production-inventory system with stationary demand and presented a closed-form expression for computing the base-stock level in the context of expected steady-state holding and backorder costs minimisation. Gallego & Özer (2001) showed that the state-dependent (s, S) policy and base-stock policy are optimal for stationary stochastic inventory systems with and without fixed ordering costs. Iida & Zipkin (2006) developed a functional approximation and a simulation-based method for approximating the base-stock policy. Lu et al. (2006) developed easy-to-compute bounds to the optimal base-stock level, which generalised and improved existing bounds in the literature. These bounds were further used to construct near-optimal policies. Chen & Lee (2009) showed that many commonly adopted time-series processes such as the auto-regressive AR(1) model, the integrated moving average IMA(0,1,1) model, the general auto-regressive moving average ARMA model could be interpreted as special cases of the MMFE framework.

These existing studies generally assumed that forecasts represent the conditional mean of demand given all information available at the time the forecast was made. Recently, Norouzi & Uzsoy (2014) considered updates of conditional covariance of demand in addition to conditional mean. They showed that the optimal base-stock level depends on the conditional covariance, and the proposed approach yields significant cost reductions and effective decisions. Albey et al. (2015) integrated the MMFE developed by Norouzi & Uzsoy (2014) to a chance-constrained stochastic optimization model in a rolling horizon setting. Computational studies, using data from a major semiconductor manufacturer, demonstrated that considering forecast evolution in the production planning model can lead to improved performance. Additional work on the multi-echelon inventory system with MMFE was conducted by (Dong & Lee, 2003; Ziarnetzky et al., 2018).

Inventory models with MMFE generally neglected the fixed ordering cost and aimed at computing the base-stock level, with the exception of Gallego &

Özer (2001), who focused on computing (s, S) policies. This paper takes into account the fixed ordering cost and presents a solution method based on the (R, S) policy. Besides, the MMFE assumes forecast updates are stationary. Our solution method can tackle demand updates following MMFEs, as well as nonstationary stochastic distributions.

3. A stochastic dynamic programming formulation

We consider a nonstationary stochastic lot-sizing problem over a T -period planning horizon. The demand d_t in each period $t = 1, \dots, T$ is a random variable which follows a probability density function $g_{d_t}(\cdot)$ and a cumulative density function $G_{d_t}(\cdot)$. Let η_t represent the realisation of d_t , and $F_t = (\eta_1, \dots, \eta_{t-1})$ denote demand realisations up to the beginning of period t . In contrast to most exiting literature, we assume that demands in successive periods are not independently distributed, but correlated. Thus, d_t follows the conditional probability density function $g_{d_t}(\zeta_t|F_t)$, where ζ_t represents the value of random variable d_t . A full list of symbols is available in Appendix A.

For our model, we will use some basic results from multivariate analysis. Let d be a n -variate multivariate normal random variable with mean $\tilde{d}$ and covariance Σ , abbreviated $d \sim \mathcal{MVN}(\tilde{d}, \Sigma)$. Let $d = [d_1 \ d_2]^T$ be a partitioned multivariate normal random n -vector, $d_1 = [d_1 \ \dots \ d_p]^T$, $d_2 = [d_{p+1} \ \dots \ d_n]^T$, with mean $\tilde{d} = [\tilde{d}_1 \ \tilde{d}_2]^T$ and covariance $\Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}$. Then the conditional distribution of d_2 given $d_1 = \eta_1$ is normally distributed with mean $\tilde{d}_2 + \Sigma_{21}\Sigma_{11}^{-1}(\eta_1 - \tilde{d}_1)$ and variance $\Sigma_{22} - \Sigma_{21}\Sigma_{11}^{-1}\Sigma_{12}$ (Theorem 3.1).

Theorem 3.1 (Conditional distribution (see for instance Billingsley (2008))).

$$E[d_2|d_1 = \eta_1] = \tilde{d}_2 + \Sigma_{21}\Sigma_{11}^{-1}(\eta_1 - \tilde{d}_1) \quad (1)$$

$$\text{Var}(d_2|d_1 = \eta_1) = \Sigma_{22} - \Sigma_{21}\Sigma_{11}^{-1}\Sigma_{12}. \quad (2)$$

Regarding the dynamics of the inventory, we assume that replenishments Q_t are placed at the beginning of period t given the opening inventory level

I_{t-1} and demand realisations $F_t = (\eta_1, \dots, \eta_{t-1})$, and delivered instantaneously. Note that I_{t-1} represents the inventory at the beginning of period t and the inventory level at the end of period $t - 1$. The ordering cost $u(\cdot)$ comprises the fixed ordering cost K , and the unit ordering cost c , as illustrated in Eq. (3).

$$u(Q_t) = \begin{cases} K + c \cdot Q_t & Q_t > 0 \\ 0 & Q_t = 0 \end{cases} \quad (3)$$

At the end of period t , the linear holding cost h is charged on every unit carried from one period to the next, and the linear penalty cost b is occurred for each unit of unmet demand.

Given the above problem description, the objective is to schedule replenishment plans to minimize the expected total cost. In what follows, we formulate this problem as a stochastic dynamic program (Bellman, 1957) comprising the following elements.

1. **State.** The state includes the inventory level I_{t-1} and demand realisations $F_t = (\eta_1, \dots, \eta_{t-1})$ at the beginning of period t .
2. **Action.** An action means to schedule a replenishment with quantity Q_t at the beginning of period t , where $Q_t \in [0, \infty)$.
3. **Expected immediate cost.** Let $f_t(I_{t-1}, F_t; Q_t)$ denote the expected immediate cost comprising ordering, holding, and penalty costs in period t , given initial inventory level I_{t-1} and demand realisations F_t .

$$\begin{aligned} f_t(I_{t-1}, F_t; Q_t) = & u(Q_t) + h \int_{d_t} \max(I_{t-1} + Q_t - \zeta_t, 0) g_{d_t}(\zeta_t | F_t) d(\zeta_t) \\ & + b \int_{d_t} \max(\zeta_t - I_{t-1} - Q_t, 0) g_{d_t}(\zeta_t | F_t) d(\zeta_t). \end{aligned} \quad (4)$$

4. **Objective function.** Let $C_t(I_{t-1}, F_t)$ denote the expected total cost of an optimal policy over periods $t, \dots, T$ with initial inventory level I_{t-1} and up-to-date demand realisations F_t . Then, $C_t(I_{t-1}, F_t)$ can be written

as, for $t = 1, \dots, T - 1$,

$$C_t(I_{t-1}, F_t) = \min_{Q_t \geq 0} \{f_t(I_{t-1}, F_t; Q_t) + \int_{d_t} C_{t+1}(I_{t-1} + Q_t - \zeta_t, F_{t+1}) g_{d_t}(\zeta_t | F_t) d(\zeta_t)\}, \quad (5)$$

where

$$C_T(I_{T-1}, F_T) = \min_{Q_T \geq 0} \{f_T(I_{T-1}, F_T; Q_T)\} \quad (6)$$

represents the boundary condition.

4. A stochastic optimisation model under (R, S) policy

In this section, we reformulate the stochastic dynamic program presented in Section 3 as a stochastic optimization model under the (R, S) policy, since our solution method is built upon modelling techniques introduced in Rossi et al. (2015) which operates under the (R, S) policy. We then combine the model with results from multivariate stochastic process analysis. The resulting stochastic optimisation model computes (R, S) policies in a way that is similar to the independent demand case. This model will further be approximated as an MILP model by adopting an existing piecewise linear approximation strategy (Rossi et al., 2015) in the next section, which can be easily solved by using off-the-shelf software.

Under the (R, S) policy, the timing of inventory reviews R and respective order-up-to-levels S are fixed simultaneously at the beginning of the planning horizon, and actual ordering quantities are decided at the beginning of each inventory review period to reach order-up-to-levels. This policy provides effective means for reducing planning instability and coping with demand uncertainty (Bookbinder & Tan, 1988). Many effective methods have been proposed for the computation of (R, S) policy parameters under the assumption that the demand is independently distributed, such as (Bookbinder & Tan, 1988; Tarim & Kingsman, 2004; Rossi et al., 2015). However, this paper leverages existing

modelling techniques and results from multivariate stochastic process analysis for computing (R, S) policy parameters under correlated demand.

We first introduce a binary variable δ_t , $t = \{1, \dots, T\}$, which takes value 1 if a replenishment is placed at the beginning of period t and 0 otherwise. Then, the ordering cost in Eq. (3) can be rewritten as follows,

$$u(Q_t) = K\delta_t + c \cdot Q_t. \quad (7)$$

At the beginning of the planning horizon, the initial inventory level is I_0 , and we observed no demand realisations, i.e., $F_t = \emptyset$. Following the (R, S) policy, the objective is to decide the timing of replenishments $\{\delta_1, \dots, \delta_T\}$ and the corresponding order-up-to-levels $\{S_1, \dots, S_T\}$ for the entire planning horizon so that the expected total cost $C_1(I_0)$ is minimized. Thus, the problem described in Section 3 can be formulated as the following stochastic optimization model.

$$\begin{aligned} \bar{C}_1(I_0) = & \min_{\substack{\delta_1, \dots, \delta_T \\ S_1, \dots, S_T}} \int_{d_1} \cdots \int_{d_T} \sum_{t=1}^T \left(K\delta_t + c \cdot Q_t + h \max(I_{t-1} + Q_t - \zeta_t, 0) \right. \\ & \left. + b \max(\zeta_t - I_{t-1} - Q_t, 0) \right) g_{d_1}(\zeta_1 | F_1) \cdots g_{d_T}(\zeta_T | F_T) d(\zeta_1) \cdots d(\zeta_T) \end{aligned} \quad (8)$$

subject to, $t = 1, \dots, T$

$$Q_t = (S_t - I_{t-1})\delta_t \quad (9)$$

$$I_t = I_0 + \sum_{i=1}^t (Q_i - \zeta_i) \quad (10)$$

$$Q_t, S_t \geq 0, I_t \in \mathcal{R}, \delta_t \in \{0, 1\} \quad (11)$$

Figure 1: A stochastic optimization model of (R, S) policy under correlated demand

The objective function (8) fixes the timing of replenishments and the associated order-up-to-levels once and for all at the beginning of the planning horizon. Constraints (9) state that the ordering quantity Q_t is equal to order-up-to-level, minus the opening inventory level at the beginning of period t if an

order is placed, and 0 otherwise. Constraints (10) are the inventory conservation constraints. The inventory level at the end of period t must be equal to the inventory level at the beginning of this period, plus all received replenishments, minus demand realisations up to period t . Constraints (11) specify domains of the order quantity, order-up-to level, inventory level, and binary variable δ_t .

Lemma 4.1 (Law of total expectation (Weiss, 2006)). *If X is an integral random variable (i.e. $E[X] < \infty$) and Y is any random variable, not necessarily integral, on the same probability space, then*

$$E[X] = E(E(X|Y)) \quad (12)$$

i.e., the expected value of the conditional expected value of X given Y is the same as the expected value of X .

By applying Lemma 4.1, the objective function (8) in Fig. 1 can be simplified

$$\begin{aligned} SPC_1(I_0) = & \min_{\substack{\delta_1, \dots, \delta_T \\ S_1, \dots, S_T}} \int_{d_1} \cdots \int_{d_T} \sum_{t=1}^T \left(K\delta_t + c \cdot Q_t + h \max(I_{t-1} + Q_t - \zeta_t, 0) \right. \\ & \left. + b \max(\zeta_t - I_{t-1} - Q_t, 0) \right) g_{d_1}(\zeta_1) \cdots g_{d_T}(\zeta_T) d(\zeta_1) \cdots d(\zeta_T). \end{aligned} \quad (13)$$

Existing studies modelled correlated demand as a time-series process (Fotopoulos et al., 1988; Graves, 1999) and MMFE (Dong & Lee, 2003; Albey et al., 2015; Ziarnetzky et al., 2018). These studies generally update future demand distributions with conditional means and covariances based on up-to-date demand realisations. However, Eq. (13) shows that the future demand distributions may be updated by leveraging unconditional means and covariances, i.e. without considering demand realisations. Therefore, under Lemma 4.1, *the nonstationary stochastic inventory problems under correlated demand can be tackled in a way that is similar to the independent demand case.*

5. A new approach for nonstationary inventory problems under receding horizon control

This section first formulates the stochastic optimization model presented in Section 4 as a mixed integer linear programming (MILP) model (Section 5.1) by using the modelling technique introduced in (Rossi et al., 2014, 2015). This model is then implemented under the receding horizon control in Section 5.2. The solution method presented is further adapted to tackle demand updates following MMFEs and time-series processes in Section 5.3.

5.1. An MILP model for computing (R, S) policies with correlated demand

This section formulates the stochastic optimization model in Fig. 1 as an MILP model for approximating (R, S) policies under correlated demand, which is built upon the modelling techniques introduced in Rossi et al. (2015).

Consider a random variable ω and a scalar variable x . The first order loss function is defined as $L(x, \omega) = E[\max(\omega - x, 0)]$, where E denotes the expected value with respect to the random variable ω . The complementary first order loss function is defined as $\hat{L}(x, \omega) = E[\max(x - \omega, 0)]$. Like Rossi et al. (2015), we model non-linear holding and penalty costs by means of this function.

Let d_{jt} represent the convolution $d_j + \dots + d_t$. Since d_t follows the multivariate normal distribution $d_t \sim \mathcal{MVN}(\tilde{d}_t, \Sigma)$, the demand convolution d_{jt} follows the normal distribution $d_{jt} \sim \mathcal{N}(\tilde{d}_{jt}, \sigma_{jt}^2)$, where $\tilde{d}_{jt} = \tilde{d}_j + \dots + \tilde{d}_t$ and $\sigma_{jt}^2 = \mathbb{1}^T \sum_{j,t} \mathbb{1}$, and $\tilde{d}$ denotes the expected value of d (Taboga, 2012).

Let ζ_{it} denote the value of random variable d_{it} , for $t = \{i, \dots, j\}$. Consider a replenishment cycle $i, \dots, t$, in which the only replenishment is placed at the beginning of period i . The inventory level at the end of period t must be equal to the order-up-to-level at the beginning of period i , minus the demand convolution over periods $i, \dots, t$, i.e. $I_t = S_i - \zeta_{it}$. Therefore, the expected on-hand stock at the end of period t in Eq. (4) can be reformulated in the form of the complementary of the first order loss function,

$$\hat{L}(S_i, d_{it}) = \int_{d_i} \dots \int_{d_t} \max(S_i - \zeta_{it}, 0) g_{d_i}(\zeta_i) \dots g_{d_t}(\zeta_t) d(\zeta_i) \dots d(\zeta_t). \quad (14)$$

And, the expected back-order at the end of period t can be reformulated in the form of first order loss function,

$$L(S_i, d_{it}) = \int_{d_i} \dots \int_{d_t} \max(\zeta_{it} - S_i, 0) g_{d_i}(\zeta_i) \dots g_{d_t}(\zeta_t) d(\zeta_i) \dots d(\zeta_t). \quad (15)$$

We introduce a binary variable P_{jt} which is set to one if the most recent replenishment up to period t was issued in period j , where $j \leq t$; if no replenishment occurs before or at period t , then we let $P_{1t} = 1$, this allows us to properly account for demand variance from the beginning of the planning horizon. We observe that if $P_{jt} = 1$, the closing inventory level of period t must be equal to the order-up-to-level of period j minus the demand convolution over periods $j, \dots, t$, i.e. $I_t = S_j - \zeta_{jt}$. Then, following Eq. (14) - (15), the expected on-hand stock and back-order of period t can be written by means of the complementary of the first order loss function and the first order loss function, $\sum_{j=1}^t \hat{L}(S_j, d_{jt}) P_{jt}$, and $\sum_{j=1}^t L(S_j, d_{jt}) P_{jt}$. Additionally, since period j must be the only most recent order received up to period t , the following constraints must be satisfied.

$$\sum_{j=1}^t P_{jt} = 1, \quad (16)$$

$$P_{jt} \geq \delta_j - \sum_{k=j+1}^t \delta_k, \quad j = 1, \dots, t. \quad (17)$$

We next employ the piecewise linear approximation technique proposed in Rossi et al. (2014, 2015) to approximate the expected on-hand stock $\hat{L}(S_j, d_{jt})$ and back-order $L(S_j, d_{jt})$. This technique requires first to partition the support Ω of d_{jt} into W disjoint subregions $\Omega_1, \dots, \Omega_W$. We define the probability mass $p_i = \Pr[d_{jt} \in \Omega_i]$, and the conditional expectation $E[d_{jt}|\Omega_i]$ with associated region Ω_i , for $i = 1, \dots, W$. Let $\hat{L}_{lb}(S_j, d_{jt})$ represent the lower bound of the expected on-hand stocks. Based on Jensen's inequality (Birge & Louveaux (2011), p. 120),

$$\hat{L}(S_j, d_{jt}) \geq \hat{L}_{lb}(S_j, d_{jt}) = \sum_{i=1}^W p_i \max(S_j - E[d_{jt}|\Omega_i], 0). \quad (18)$$

Since $L(S_j, d_{jt}) = \hat{L}(S_j, d_{jt}) - (S_j - \tilde{d}_{jt})$ (Snyder & Shen, 2019) and $S_j - \tilde{d}_{jt} = \tilde{I}_t$, the i^{th} segment of the lower bound of the back-order $L_{lb}(S_j, d_{jt})$ can be written as

$$L_{lb}^i(S_j, d_{jt}) \geq -\tilde{I}_t + \sum_{k=1}^i p_k S_j - \sum_{k=1}^i E[d_{jt}|\Omega_k] \quad E[d_{jt}|\Omega_i] \leq S_j \leq E[d_{jt}|\Omega_{i+1}]. \quad (19)$$

Let e_W^{jt} denote the approximation error. The Edmundson-Madansky upper bounds are obtained by shifting the lower bounds up by e_W^{jt} . Let $\tilde{H}_t \geq 0$ and $\tilde{B}_t \geq 0$ denote the upper bounds to the expected on-hand and back-order stocks at the end of period t , then they are formulated as follows,

$$\tilde{H}_t \geq \sum_{j=1}^t S_j P_{jt} \sum_{k=1}^i p_k + \sum_{j=1}^t \left(e_W^{jt} - \sum_{k=1}^i p_k E[d_{jt}|\Omega_i] \right) P_{jt}, \quad (20)$$

$$\tilde{B}_t \geq -\tilde{I}_t + \sum_{j=1}^t S_j P_{jt} \sum_{k=1}^i p_k + \sum_{j=1}^t \left(e_W^{jt} - \sum_{k=1}^i p_k E[d_{jt}|\Omega_i] \right) P_{jt}, \quad (21)$$

$t = 1, \dots, T$, and $i = 1, \dots, W$. Note that $\sum_{j=1}^t S_j P_{jt} = \tilde{I}_t + \sum_{j=1}^t \tilde{d}_{jt} P_{jt}$.

Therefore, the stochastic optimization model in Fig. 1 for tackling the non-stationary stochastic inventory problem under correlated demand can be formulated as an MILP model in Fig. 2.

The objective (22) decides the timing and order-up-to-level of replenishments so as to minimize the expected total cost comprising ordering, holding, and penalty costs given the initial inventory level I_0 . Constraints (23) ensure the non-negativity of replenishments. Constraints (24) are indicator constraints (Belotti et al., 2016) capturing the reorder condition. Constraints (25) indicate the most recent replenishment before period t was issued in period j . Constraints (26) uniquely define in which the most recent replenishment prior to t took place. Constraints (27) - (28) are approximations of expected holding and penalty costs at the end of period t utilising the first order loss function and its complementary function. Constraints (29) - (30) identify binary variables.

The MILP model in Fig. 2 can be easily implemented and solved using off-the-shelf software. The timing of inventory reviews are obtained from δ_t and

$$\min_{\delta_t} -cI_0 + c \sum_{t=1}^T \tilde{d} + \sum_{t=1}^T (K\delta_t + h\tilde{H}_t + b\tilde{B}_t) + c\tilde{H}_T \quad (22)$$

Subject to, for $t = 1, \dots, T$, $j = 1, \dots, t$, and $i = 1, \dots, W$,

$$\tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1} \geq 0 \quad (23)$$

$$\delta_t = 0 \rightarrow \tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1} = 0 \quad (24)$$

$$\sum_{j=1}^t P_{jt} = 1, \quad (25)$$

$$P_{jt} \geq \delta_j - \sum_{k=j+1}^T \delta_k, \quad (26)$$

$$\tilde{H}_t \geq (\tilde{I}_t + \sum_{j=1}^t \tilde{d}_{jt} P_{jt}) \sum_{k=1}^i p_k + \sum_{j=1}^t \left(e_W^{jt} - \sum_{k=1}^i p_k \mathbb{E}[d_{jt} | \Omega_i] \right) P_{jt}, \quad (27)$$

$$\tilde{B}_t \geq -\tilde{I}_t + (\tilde{I}_t + \sum_{j=1}^t \tilde{d}_{jt} P_{jt}) \sum_{k=1}^i p_k + \sum_{j=1}^t \left(e_W^{jt} - \sum_{k=1}^i p_k \mathbb{E}[d_{jt} | \Omega_i] \right) P_{jt}, \quad (28)$$

$$\delta_t \in \{0, 1\} \quad (29)$$

$$P_{jt} \in \{0, 1\} \quad (30)$$

Figure 2: An MILP model for computing (R, S) policies with correlated demands

the corresponding order-up-to-levels S_t are obtained from $\tilde{I}_t + \tilde{d}_t$.

For the special case in which demand follows a standard normal distribution, the piecewise linear approximation parameters p_i , $\mathbb{E}[d_{jt} | \Omega_i]$, and e_W^{jt} are provided in Rossi et al. (2014). These parameters can be applied to general normal distributions by using the standardisation formula $\hat{L}(S_j, d_{jt}) = \sigma_{jt} \hat{L}\left(\frac{S_j - \tilde{d}_{jt}}{\sigma_{jt}}, Z\right)$ in Rossi et al. (2014), Lemma 7, where Z is a standard normal random variable. Note that, since in our case demands are correlated, $d_{jt} \sim \mathcal{N}(\tilde{d}_{jt}, \sigma_{jt}^2)$, where $\tilde{d}_{jt} = \tilde{d}_j + \dots + \tilde{d}_t$ and $\sigma_{jt}^2 = \mathbb{1}^T \sum_{j_t} \mathbb{1}$.

5.2. An MILP model under receding horizon control

This section implements the MILP model in Fig. 2 under a receding horizon control framework, which is widely used in the inventory control literature (Kilic & Tarim, 2011; Dural-Selcuk et al., 2020).

Our algorithm proceeds as follows. At the beginning of each period k , the inventory level I_{k-1} and up-to-date demand realisations F_k are observed, based on which demand distributions for the remaining periods of the planning horizon $k, k+1, \dots, T$ are updated using Theorem 3.1. Then, inventory review periods and the corresponding order-up-to levels are obtained by solving the $T - k + 1$ periods problem using the MILP model in Fig. 2. Next, the replenishment plan for the current period k is put into action, i.e., an order is placed if necessary. At the end of this period, the demand realisation η_k is observed, and the expected total cost over periods $1, 2, \dots, k$ is calculated. This procedure repeats until the end of the planning horizon. We simulate this process until the stopping criterion of a maximum estimation error of the expected total cost under a given confidence probability is satisfied.

The MILP model under the receding horizon control is structured as follows (Algorithm 1).

Update of opening inventory level and demand forecasts (line 5-6).

We update the inventory level at the beginning of period k , I_{k-1} , and demand forecasts $E[d_t|F_k]$ and $\text{Cov}(d_t, d_{t+i}|F_k)$ for $t = k, \dots, T$ and $i = 0, \dots, T - k$ by using the conditional distribution, Theorem 3.1.

Computation of optimal (R, S) policies (line 7). We obtain current replenishment decisions δ_k and Q_k by solving the MILP model (Fig. 2) over periods $k, \dots, T$.

Implementation of imminent replenishment decision (line 8). The replenishment decision δ_k and Q_k in period k is put into action.

Calculation of total cost (line 9). We calculate the total cost over periods $1, \dots, k$.

Receding horizon control differs from rolling horizon control, in which a fixed length of time window rolls many times over the planning horizon. Rolling

Input : costs (*orderingcost*, *holdingcost*, *penaltycost*),
expecteddemand, *covariance*, *coefficientofcorrelation*,
initialStock, *errorThreshold*, *confidenceProbability*

Output: *totalCost*

```

1 while error > errorThreshold do
2   generate demand realisations  $\eta = (\eta_1, \dots, \eta_T)$ ;
3   expectedTotalCost = 0;
4   for  $k = 1$  to  $T$  do
5     update opening inventory level  $I_{k-1} = I_{k-2} + Q_{k-1} - \eta_{k-1}$ ;
6     update demand forecasts  $E[d_k|F_k]$  and  $\text{Cov}(d_k, d_{k+i}|F_k)$ ,
        $i = 0, \dots, T - k$ ;
7     solve the MILP model in Fig. 2 with  $I_{k-1}$ ,  $E[d_k|F_k]$  and
        $\text{Cov}(d_k, d_{k+i}|F_k)$ ;
8     implement imminent  $\delta_k$  and  $Q_k$ ;
9     calculate up-to-date costs  $\text{totalCost} += f_k(I_{k-1}, F_k; Q_k)$ ;
10  end
11  calculate error;
12 end
13 return totalCost

```

Algorithm 1: An MILP model under receding horizon control

horizon control is more suitable for infinite or long-horizon planning problems. Besides, the length of the time window may significantly affect the optimal replenishment plans, as pointed out in Bookbinder & Tan (1988). Receding horizon control sets the length of the time window equal to the number of remaining periods in the planning horizon, thus capturing end-of-horizon effects that affect products with short life cycles (Dural-Selcuk et al., 2020).

5.3. Extensions to the Martingale Model of Forecast Evolution

The Martingale Model of Forecast Evolution (MMFE) proposed by (Graves et al., 1986; Heath & Jackson, 1994) is a powerful framework that can capture

both time-series models of prediction and judgemental forecasts. In this section, we demonstrate that our model presented in section 5.2 can be adapted to capture the MMFE.

Under the MMFE framework, demand forecasts are assumed to be available for a certain number of periods in the future. Let H be the forecast horizon. At the end of period s , demand forecasts for the following H periods are generated. Let $D_{s,t}$, $s \leq t \leq s + H$, represent the demand forecast made in period s for period t . $D_{s,s} = D_s$ denotes the demand realisation in period s , since the forecasts are made after the actual demand of period s is revealed. The demand forecast $D_{s,t}$, $t > s + H$, is set to a constant μ . Note that the MMFE requires the demand to be stationary (Heath & Jackson, 1994). However, our model has the advantage of being able to tackle nonstationary demands. Heath & Jackson (1994) developed two classes of models for the behaviour of forecast updates: the additive model and the multiplicative model. For exposition simplicity, we mainly focus on the additive model in the rest of this section. The multiplicative model is discussed in (Heath & Jackson, 1994; Toktay & Wein, 2001; Iida & Zipkin, 2006; Norouzi & Uzsoy, 2014).

Let $\epsilon_{s,t} = D_{s,t} - D_{s-1,t}$ be a random variable which represents the forecast update made at the end of period s for period t . $\epsilon_{s,t} = 0$ for $t > s + H$. Let $\epsilon_s = (\epsilon_{s,s}, \epsilon_{s,s+1}, \dots, \epsilon_{s,s+H})$ be the forecast update vector generated at the end of period s . $\epsilon_{s,s}$ represents the final update made to period s , and $\epsilon_{s,s+H}$ denotes the first forecast update made to period $s+H$. Heath & Jackson (1994) assumed the update vectors ϵ_s are independent and identically distributed multivariate normal random vectors with mean 0 and covariance matrix $\Sigma = \sigma_{i,j}$, where $\sigma_{i,j}$ represents the covariance of $\epsilon_{s,s+i}$ and $\epsilon_{s,s+j}$, $i, j = 0, 1, \dots, H$.

Since the MMFE requires the demand to be stationary (Heath & Jackson, 1994; Dong & Lee, 2003; Norouzi & Uzsoy, 2014), the unconditional mean of the demand process is a constant μ , and the conditional mean for period t given

information available in period s is

$$D_t = \mu + \sum_{j=0}^H \epsilon_{t-H+j,t}. \quad (31)$$

Toktay & Wein (2001) developed the following unconditional covariance between D_t and D_{t+i} ,

$$\text{Cov}(D_t, D_{t+i}) = \sum_{j=0}^H \sigma_{j,i+j}, \quad i = 0, 1, \dots, H. \quad (32)$$

The unconditional covariance, Eq. (32), ignores the available demand information that forecast updates provided. Norouzi & Uzsoy (2014) incorporated this information and proposed the following conditional covariance between D_t and D_{t+i} given all information available at the end of period s ,

$$\text{Cov}(D_t, D_{t+i}|F_s) = \sum_{j=1}^{t-s} \sigma_{t-s-j,t+i-s-j}, \quad 1 \leq t-s \leq H-i+1. \quad (33)$$

When $s = t$, D_t is totally revealed; all covariance matrix elements $\text{Cov}(D_t, D_{t+i}|F_s) = 0$.

The MMFE differs from the multivariate normal distribution because of the length of the forecast horizon. Given the information available at the beginning of a certain period, the MMFE framework updates demand forecasts for the next H periods. In contrast, the multivariate normal distribution updates demand forecasts for the remaining periods of the planning horizon. Recall that our solution method consists of an MILP model (Fig. 2) and its implementation under the receding horizon control framework (Algorithm 1). The MILP model tackles the nonstationary stochastic inventory problem with correlated demand in a way which is similar to the independent demand case, which requires the unconditional mean and covariance of demands. Given the information available at the beginning of the planning horizon, the unconditional mean of MMFE is a constant μ , and its unconditional covariance for periods $1, \dots, H+1$ is updated via Eq. (32). Thus, our MILP model can be adapted to solve demands which follow MMFEs. Under a receding horizon control framework, given the information available at the beginning of period t , the conditional expected

demand and conditional covariance for periods $t, \dots, t + H$ are updated based on Eq. (31) and Eq. (33). Therefore, our solution method presented in Section 5.2 can be immediately adapted to tackle stochastic inventory problems with demand updates following MMFEs.

Most commonly used time-series models such as the AR, MA, and ARMA processes are special cases of the MMFE (Dong & Lee, 2003; Iida & Zipkin, 2006; Norouzi & Uzsoy, 2014). Therefore, our model presented in Section 5.2 can also be adapted to tackle demands which follow all these time-series processes.

6. Computational experiments

This section presents a numerical study to investigate the cost performance of the solution method discussed in Section 5. We assume that the demand follows a multivariate normal distribution featuring correlation over different periods. Note that our approach comes with the advantage of tackling demands following a collection of time-series processes and the MMFE built upon the multivariate normal distribution; without loss of generality, we thus only consider demands following multivariate normal distributions in this section. Numerical experiments are conducted using IBM ILOG CPLEX Optimization Studio 12.7 and Eclipse 4.7.3 on a 3.2GHz Intel(R) Core(TM) with 8GB of RAM, and MATLAB R2019a on a 3.6GHz Intel(R) Xeon(R) with 48GB of RAM.

We consider nine demand patterns: two life cycle patterns (LCY1 and LCY2), two sinusoidal patterns (SIN1 and SIN2), a stationary pattern (STA), and four empirical patterns (EMP1, ..., EMP4). These demands are normally distributed with means presented in Appendix B and coefficients of variation $c_v \in \{0.1, 0.2\}$, where $\sigma_{d_t} = c_v \cdot \tilde{d}_t$. The fixed ordering cost K , proportional ordering cost c and penalty cost b range in $\{150, 300\}$, $\{0, 1\}$, and $\{5, 10, 20\}$, respectively. The proportional holding cost is $h = 1$, and the coefficient of correlation $\rho \in \{-0.75, -0.25, 0.25, 0.75\}$.¹

¹A negative coefficient of correlation describes the extent to which two variables move

We utilise the SDP model discussed in Section 3 as a benchmark. Note that the SDP is pseudo-polynomial; therefore, an increase in the average value of the demand or its standard deviation leads to a dramatic increase in the state space boundaries and computational times. A full factorial experimental design is computational impractical. We therefore adopt Latin Hypercube Sampling (McKay et al., 2000) to obtain ten instances for each demand pattern. In total, we consider 90 instances in our computational study. These instances are presented in Appendix B.

We first implement the SDP model in Matlab 2019a. Since the demand follows a multivariate normal distribution, we adopt a discretisation step size of 0.5 and a continuity correction factor of 0.25 for demand patterns LCY1, LCY1, SIN1, SIN2, and STA to guarantee the model accuracy, and a discretisation step size of 1 and a continuity correction factor of 0.5 for EMP demand patterns (EMP1, EMP2, EMP3, EMP4) to ensure a reasonable computational time. Given a set of demand realisations, future demand distributions are updated based on the conditional distribution (Theorem 3.1). Since demands are correlated over different periods, it is computationally impractical to find the optimal solution using SDP (Levi et al., 2008; Nasr & Elshar, 2018). Therefore, we propose a special structure $\sigma_{ij}^2 = \rho^{|j-i|} \sigma_i \sigma_j$ to construct the covariance matrix. Accordingly, the computation of the conditional mean only requires knowledge of the realised demand in the previous period, which in return reduces the size of the state space.

We then implement the MILP model (Fig. 2) in the IBM ILOG CPLEX Optimization Studio 12.7 to obtain (R, S) policies, and simulate these policies 1 million times to obtain the expected total costs. Since these policies are ob-

in opposite directions; it captures situations such as pre- and post-holiday season demand, e.g. customers tend to purchase more during Christmas sales, summer sales, and market campaigns, and less in the months that immediately follow such seasons. A positive coefficient of correlation describes the extent to which two variables move in the same direction. It captures the situations such as new products launch, social network marketing, and stock market, e.g. new products' sales surge as positive feedback emerges.

tained based on the unconditional distribution and this model doesn't take into account demand realizations, we use "static method" to denote this model. We further implement the MILP model under the receding horizon control framework (Algorithm 1) in Eclipse 4.7.3, which updates the (R, S) policies at the beginning of each time period based on up-to-date demand realisations. We use "dynamic method" to denote Algorithm 1 presented in Section 5.2. Since we operate under the assumption of normality, our models can be readily linearized by using the piecewise linearization parameters available in Rossi et al. (2014). Specifically, we employ eleven segments in the piecewise-linear approximations of the expected on-hand stocks and back-orders. For the dynamic method, we adopt the stopping criterion of a maximum estimation error of 0.02% of the expected total cost with 98% confidence probability. We compare the expected total costs obtained via the static method and dynamic method against that obtained via SDP.

We present box plots of the optimality gaps of the static method and dynamic method in Fig. 3. Note that the y -axis is displayed in logarithmic scale. The average optimality gap of static method and dynamic method are 0.68% and 0.16%, and their worst-case performance are 6.86% and 3.02%, respectively. These differences are expected due to the approximate nature of both these methods, as originally reported in Dural-Selcuk et al. (2020) for the case of uncorrelated demand. We observe that a receding horizon control framework can provide tighter optimality gaps and reduce their dispersion.

We then present optimality gaps of the dynamic method for different pivoting parameters in Table Appendix C. The results presented confirm that our approach is robust and consistently leads to near-optimal solutions and low dispersion under all pivoting parameters.

We further investigate the cost of ignoring demand correlation. We compare the optimality gaps when demand correlation is considered and ignored in Fig 4. When demand correlation is ignored, the system is controlled via an (s, S) policy computed by setting $\rho = 0$. Note that the y -axis is displayed in logarithmic scale. We observe that when the planning horizon is short and the demand cor-




Figure 3: Optimality gaps % of the static method and the dynamic method

relation is weak, as one may expect, the difference between considering demand correlation and ignoring demand correlation becomes negligible. We therefore consider a longer planning horizon example with a higher correlation coefficient.

We consider a 29-period example whose expected demands are shown in Fig. 5; this instance closely resembles empirical demand patterns presented in Kurawarwala & Matsuo (1996). The fixed and proportional ordering costs are 300 and 0; the proportional holding and penalty costs are 1 and 5, respectively. Coefficient of variation is $cv = 0.1$; and coefficient of correlation is $\rho = 0.9$. We first solve this 29-period problem by using the dynamic method proposed in Algorithm 1, and obtain the expected total cost 6314.64. We then solve the problem by ignoring the correlation, and obtain the expected total cost 6813.58. The cost difference observed is therefore 7.90%. This demonstrates that ignoring demand correlation may lead to a substantially more expensive control action.

7. Conclusion

This paper considered the single-item single-stocking location nonstationary stochastic inventory problem under correlated demand. We presented the




Figure 4: Optimality gap % for different correlation coefficients




Figure 5: Expected demand of the 29-period example ($K = 300$, $c = 0$, $h = 1$, $b = 5$, $cv = 0.1$, and $\rho = 0.9$), for which the cost difference observed when correlation is ignored is 7.90%.

first mathematical programming-based solution method to model and solve the problem by combining an existing piecewise linear approximation strategy with results from multivariate stochastic process analysis in the context of horizon

control framework. The resulting MILP model can be easily solved by using off-the-shelf software. We also showed that our solution method can be adapted to tackle demand under various assumptions: the multivariate normal distribution, a collection of time-series processes, and the MMFE.

We conducted a numerical study including 90 8-period instances extrapolated from a test bed in the literature by means of a stratified sampling strategy. We compared the expected total costs obtained via our solution method with the optimal ones obtained via SDP. Computational experiments showed that our mathematical programming-based solution method yields a mean optimality gap of 0.16% and an IQR of 0.07%, which demonstrate that our approach leads to near-optimal solutions and low dispersions. Finally, we conducted a 29-period experiment which showed that under highly correlated demand, ignoring demand correlation may lead to a substantially more expensive control action.

References

- Albey, E., Norouzi, A., Kempf, K. G., & Uzsoy, R. (2015). Demand modeling with forecast evolution: an application to production planning. *IEEE Transactions on Semiconductor Manufacturing*, 28, 374–384. doi:10.1109/TSM.2015.2453792.
- Bellman, R. (1957). Dynamic programming. *Princeton University Press*, .
- Belotti, P., Bonami, P., Fischetti, M., Lodi, A., Monaci, M., Nogales-Gómez, A., & Salvagnin, D. (2016). On handling indicator constraints in mixed integer programming. *Computational Optimization and Applications*, 65, 545–566. doi:10.1007/s10589-016-9847-8.
- Beyer, D., & Sethi, S. P. (1997). Average cost optimality in inventory models with Markovian demands. *Journal of Optimization Theory and Applications*, 92, 497–526. doi:10.1007/0-387-25477-3_1.
- Beyer, D., Sethi, S. P., & Taksar, M. (1998). Inventory models with Markovian

- demands and cost functions of polynomial growth. *Journal of Optimization Theory and Applications*, 98, 281–323. doi:10.1023/A:1022633400174.
- Billingsley, P. (2008). *Probability and measure*. John Wiley & Sons.
- Birge, J. R., & Louveaux, F. (2011). *Introduction to stochastic programming*. Springer Science & Business Media.
- Bookbinder, J. H., & Tan, J.-Y. (1988). Strategies for the probabilistic lot-sizing problem with service-level constraints. *Management Science*, 34, 1096–1108. doi:10.1287/mnsc.34.9.1096.
- Chen, L., & Lee, H. L. (2009). Information sharing and order variability control under a generalized demand model. *Management Science*, 55, 781–797. doi:10.1287/mnsc.1080.0983.
- Cheng, F., & Sethi, S. P. (1999). Optimality of state-dependent (s, S) policies in inventory models with markov-modulated demand and lost sales. *Production and Operations Management*, 8, 183–192. doi:10.1111/j.1937-5956.1999.tb00369.x.
- Dong, L., & Lee, H. L. (2003). Optimal policies and approximations for a serial multiechelon inventory system with time-correlated demand. *Operations Research*, 51, 969–980. doi:10.1287/opre.51.6.969.24920.
- 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*, 97, 102091. doi:10.1016/j.omega.2019.07.007.
- Fotopoulos, S., Wang, M.-C., & Rao, S. S. (1988). Safety stock determination with correlated demands and arbitrary lead times. *European Journal of Operational Research*, 35, 172–181. doi:10.1016/0377-2217(88)90026-4.
- Gallego, G., & Özer, Ö. (2001). Integrating replenishment decisions with advance demand information. *Management Science*, 47, 1344–1360. doi:10.1287/mnsc.47.10.1344.10261.

- Graves, S. C. (1999). A single-item inventory model for a nonstationary demand process. *Manufacturing & Service Operations Management*, 1, 50–61. doi:10.1287/msom.1.1.50.
- Graves, S. C., Meal, H. C., Dasu, S., & Qui, Y. (1986). Two-stage production planning in a dynamic environment. In *Multi-stage production planning and inventory control* (pp. 9–43). Springer.
- Heath, D. C., & Jackson, P. L. (1994). Modeling the evolution of demand forecasts with application to safety stock analysis in production/distribution systems. *IIE transactions*, 26, 17–30. doi:10.1080/07408179408966604.
- Hu, J., Zhang, C., & Zhu, C. (2016). (s, S) inventory systems with correlated demands. *INFORMS Journal on Computing*, 28, 603–611. doi:10.1504/ijbsr.2007.015372.
- Iida, T., & Zipkin, P. H. (2006). Approximate solutions of a dynamic forecast-inventory model. *Manufacturing & Service Operations Management*, 8, 407–425. doi:10.1287/msom.1060.0116.
- Johnson, G. D., & Thompson, H. E. (1975). Optimality of myopic inventory policies for certain dependent demand processes. *Management Science*, 21, 1303–1307. doi:10.1287/mnsc.21.11.1303.
- Kilic, O. A., & Tarim, S. A. (2011). An investigation of setup instability in non-stationary stochastic inventory systems. *International Journal of Production Economics*, 133, 286–292. doi:10.1016/j.ijpe.2010.04.021.
- Kurawarwala, A. A., & Matsuo, H. (1996). Forecasting and inventory management of short life-cycle products. *Operations Research*, 44, 131–150.
- Levi, R., Roundy, R. O., Shmoys, D. B., & Truong, V. A. (2008). Approximation algorithms for capacitated stochastic inventory control models. *Operations Research*, 56, 1184–1199. doi:10.1287/opre.1080.0580.

- Lian, Z., Liu, X., & Zhao, N. (2009). A perishable inventory model with Markovian renewal demands. *International Journal of Production Economics*, 121, 176–182. doi:10.1016/j.ijpe.2009.04.026.
- Lu, X., Song, J.-S., & Regan, A. (2006). Inventory planning with forecast updates: Approximate solutions and cost error bounds. *Operations Research*, 54, 1079–1097. doi:10.1287/opre.1060.0338.
- McKay, M. D., Beckman, R. J., & Conover, W. J. (2000). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 42, 55–61. doi:10.1080/00401706.2000.10485979.
- Nasr, W. W., & Elshar, I. J. (2018). Continuous inventory control with stochastic and non-stationary markovian demand. *European Journal of Operational Research*, 270, 198–217. doi:10.1016/j.ejor.2018.03.023.
- Neale, J. J., & Willems, S. P. (2009). Managing inventory in supply chains with nonstationary demand. *Interfaces*, 39, 388–399. doi:10.1287/inte.1090.0442.
- Norouzi, A., & Uzsoy, R. (2014). Modeling the evolution of dependency between demands, with application to inventory planning. *IIE Transactions*, 46, 55–66. doi:10.1080/0740817X.2013.803637.
- Ray, W. (1981). Computation of reorder levels when the demands are correlated and the lead time random. *Journal of the Operational Research Society*, 32, 27–34. doi:10.2307/2581466.
- Rossi, R., Kilic, O. A., & Tarim, S. A. (2015). Piecewise linear approximations for the static–dynamic uncertainty strategy in stochastic lot-sizing. *Omega*, 50, 126–140. doi:10.1016/j.omega.2014.08.003.
- Rossi, R., Tarim, S. A., Prestwich, S., & Hnich, B. (2014). Piecewise linear lower and upper bounds for the standard normal first order loss function. *Applied*

- Mathematics and Computation*, 231, 489–502. doi:10.1016/j.amc.2014.01.019.
- Scarf, H. (1960). Optimality of (s, S) policies in the dynamic inventory problem. In *Mathematical Methods in the Social Science* (pp. 196–202). Stanford University Press.
- Sethi, S. P., & Cheng, F. (1997). Optimality of (s, S) policies in inventory models with Markovian demand. *Operations Research*, 45, 931–939. doi:10.2139/ssrn.1126048.
- Snyder, L. V., & Shen, Z.-J. M. (2019). *Fundamentals of supply chain theory*. John Wiley & Sons.
- Song, J.-S., & Zipkin, P. (1993). Inventory control in a fluctuating demand environment. *Operations Research*, 41, 351–370. doi:10.1287/opre.41.2.351.
- Taboga, M. (2012). *Lectures on probability theory and mathematical statistics*. CreateSpace Independent Publishing Platform.
- Tarim, S. A., & Kingsman, B. G. (2004). The stochastic dynamic production/inventory lot-sizing problem with service-level constraints. *International Journal of Production Economics*, 88, 105–119. doi:10.1016/S0925-5273(03)00182-8.
- Toktay, L. B., & Wein, L. M. (2001). Analysis of a forecasting-production-inventory system with stationary demand. *Management Science*, 47, 1268–1281. doi:10.1287/mnsc.47.9.1268.9787.
- Weiss, N. A. (2006). *A course in probability*. Addison-Wesley.
- Ziarnetzky, T., Mönch, L., & Uzsoy, R. (2018). Rolling horizon, multi-product production planning with chance constraints and forecast evolution for wafer fabs. *International Journal of Production Research*, 56, 6112–6134. doi:10.1080/00207543.2018.1478461.

Appendix

Please see the supplementary materials.