MILP-based model for approximating non-stationary (R, S) policies with correlated demands

Mengyuan Xiang, Roberto Rossi, Belen Martin-Barragan
Business School, University of Edinburgh
{mengyuan.xiang,roberto.rossi,belen.martin}@ed.ac.uk

S.Armagan Tarim
Department of Management, Cankaya University
at@cankaya.edu.tr

Abstract

This paper addresses the single-item single-stock location stochastic lot-sizing problem under (R, S) policy. We assume demands in different periods are dependent. We present a mixed integer linear programming (MILP) model for computing optimal (R, S) policy parameters, which is built upon the conditional distribution. Our model can be extended to cover time-series-based demand processes as well. Our computational experiments demonstrate the effectiveness and versatility of this model.

1 Introduction

Since [14] proved the optimality of (s, S) policies for a class of dynamic inventory models, a sizeable literature has been performed for computing the optimal policy parameters (see, for example, [20, 1, 6]). However, as pointed out in [19], although the (s, S) policy is cost-optimal, it performs poorly in terms of "nervousness", i.e. lack of planning stability. In this regard, the (R, S) policy provides an effective means of dampening the planning instability and coping with demand uncertainty. Under this policy, both inventory reviews R and associated order-up-to-levels S are fixed at the beginning of the planning horizon, while actual order quantities are decided upon only after demand has been observed.

In the seminal work, [3] proposed a two-stage deterministic equivalent heuristic which fixes replenishment periods first and then determines order quantities under the independent demand assumption. [17] presented a mixed integer programming (MIP) model that determines both timing and quantity of orders simultaneously without addressing computational performance. Under the independent demand assumption, [19, 12, 16, 17, 18, 11] proposed efficient solution methods.

In the literature, most inventory models assume that demands over different time periods are independent and identically distributed. Recently, a few studies on inventory theory with correlated demands have been emerged. They either focused on (s, S) policy (see [9, 15, 5, 4]) or measured the performance of the inventory system with specific demand patterns ([9, 7, 10, 8]). However, none of them studied the (R, S) policy with correlated demands, which motivates our work in developing an efficient method for computing (R, S) policies.

In this paper, we present an MILP-based model for approximating the (R, S) policies with correlated demand. Our model can cover a series of time-based demand process, such as the autoregressive process (AR), the moving-average process (MA), the autoregressive moving average process (ARMA), the autoregressive conditional heteroskedasticity process (ARCH). Preliminary computational experiments demonstrate that optimality gaps of our model are tighter than existing algorithms, and computational times of model are reasonable. Our model can be accommodated to approximate (s, S), and (R, Q) policies.

2 Problem description

Let a random vector $\mathbf{d} = [\mathbf{d_1}, \ldots, \mathbf{d_n}]^T$ represents stochastic demand over the planning horizon, which follows the multivariate distribution f with cumulative distribution function $F : \mathcal{R}^n \to \mathcal{R}$. Let $\tilde{\mathbf{d}}$ be the mean of demand vector $\mathbf{d}$, and $\mathbf{\Sigma}$ be the variance-covariance matrix, we require that $\mathbf{\Sigma}$ is symmetric positive definite.

Lemma 1 (Conditional distribution) *Let* $\mathbf{d} == [\mathbf{d_1}, \ldots, \mathbf{d_q}, \mathbf{d_{q+1}} \ldots, \mathbf{d_n}]^T$ *denote a random vector with joint probability function* $f(d_1, \ldots, d_q, \ldots, d_n)$*, then the conditional joint probability density function of* $d_1, \ldots, d_q$ *given* $\mathbf{d_{q+1}} = d_{q+1}, \ldots, \mathbf{d_n} = d_n$ *is*

$$f_{1,\ldots,q|q+1,\ldots,n}(d_1, \ldots, d_q | \mathbf{d_{q+1}} = d_{q+1}, \ldots, \mathbf{d_n} = d_n) = \frac{f(d_1, \ldots, d_q)}{f(d_{q+1}, \ldots, d_n)} \tag{1}$$

We now consider the multivariate normal distribution (MVN). A vector-valued random variable $\mathbf{d} = [\mathbf{d_1}, \ldots, \mathbf{d_n}]^T$ is said to have a multivariate normal distribution (MVN) with mean $\tilde{\mathbf{d}} \in \mathcal{R}^n$ and covariance matrix $\mathbf{\Sigma} \in \mathcal{R}^{n \times n}$, if its probability density function is given by

$$f(d; \tilde{d}, \Sigma) = \frac{1}{(2\pi)^{\frac{1}{2}} |\Sigma|^{\frac{1}{2}}} \exp\big(- \frac{1}{2}(d - \tilde{d})^T \Sigma^{-1} (d - \tilde{d})^T \big). \tag{2}$$

Lemma 2 (Conditional distribution of MVN) *Let* $\mathbf{d} = [\mathbf{d_1}, \mathbf{d_2}]^T$ *be a partitioned multivariate normal random vector, with mean* $\tilde{\mathbf{d}} = [\tilde{\mathbf{d_1}}, \tilde{\mathbf{d_2}}]^T$ *and variance-covariance matrix*

$$\mathbf{\Sigma} = \begin{bmatrix} \mathbf{\Sigma_{11}} & \mathbf{\Sigma_{12}} \\ \mathbf{\Sigma_{21}} & \mathbf{\Sigma_{22}} \end{bmatrix}. \tag{3}$$

Then, the conditional distribution of $\mathbf{d_2}$ *given* $\mathbf{d_1} = d_1$ *is MVN, with conditional distribution* $\mathbf{d_2}|\mathbf{d_1} = d_1 \sim \mathcal{N}\big(\tilde{\mathbf{d_2}} + \mathbf{\Sigma_{21}}\mathbf{\Sigma_{11}^{-1}}(d_1 - \tilde{\mathbf{d_1}}), \mathbf{\Sigma_{22}} - \mathbf{\Sigma_{21}}\mathbf{\Sigma_{11}^{-1}}\mathbf{\Sigma_{12}}\big)$.

Example. We now demonstrate the concepts introduced on a 4-period example. $\tilde{\mathbf{d}} = [20, 40, 60, 40]$ and standard deviations $\sigma = 0.25\tilde{\mathbf{d}}$. We assume any $\mathbf{d_t}$, $t = \{2, \ldots, T\}$, is only correlated to $\mathbf{d_{t-1}}$ with correlation coefficient $\rho = 0.5$, then the variance-covariance matrix is

$$\mathbf{\Sigma} = \begin{bmatrix} 25 & 25 & 0 & 0 \\ 25 & 100 & 75 & 0 \\ 0 & 75 & 225 & 75 \\ 0 & 0 & 75 & 100 \end{bmatrix}.$$

Therefore, the conditional distribution of $\mathbf{d_t}$, for $t = \{2, \ldots, T\}$, is $\mathbf{d_2}|\mathbf{d_1} = d_1 \sim \mathcal{N}(20 + d_1, 75)$, $\mathbf{d_3}|\mathbf{d_2} = d_2 \sim \mathcal{N}(30 + \frac{3}{4}d_2, 168.75)$, and $\mathbf{d_4}|\mathbf{d_3} = d_3 \sim \mathcal{N}(20 + \frac{1}{3}d_3, 75)$.

3 Stochastic dynamic programming

We consider a single-item single-stocking location inventory management system over a T-period planning horizon. We assume that the demand d_t depends on realised information set i_{t-1} at period $t-1$; it follows the conditional distribution $f(d_t|i_{t-1})$. Let I_{t-1} denote the opening inventory level, and Q_t represent the order quantity.

At the beginning of period t, there exist ordering costs $c(\cdot)$ comprising a fixed ordering cost K, and a linear ordering cost c. At the end of period t, a linear holding cost h is charged on every unit carried from one period to the next; a linear penalty cost b is occurred for each unmet demand at the end of each time period. Then the immediate cost can be expressed as

$$f_t(i, I_{t-1}, Q_t) = c_t(i, Q_t) + \mathrm{E}[h \cdot \max(I_{t-1} + Q_t - d_t, 0) + b \cdot \max(d_t - x - Q_t, 0)|i_{t-1} = i], \tag{4}$$

where $c_t(i, Q_t)$ is defined as:

$$c_t(i, Q_t) = K \cdot \delta_t + c \cdot Q_t, \qquad \delta_t = \{0, 1\}. \tag{5}$$

Let $C_t(i, I_{t-1})$ denote the expected total cost of an optimal policy over period $t, \ldots, T$ when the observed demand information set is $i_t = i$ and the opening inventory level is I_{t-1}. We model the problem as a stochastic dynamic program ([2]) via the following functional equation,

$$C_t(i, I_{t-1}) = \min_{Q_t \geq 0} \{f_t(i, I_{t-1}, Q_t) + E[C_{t+1}(i_{t+1}, I_{t-1} + Q_t - d_t) | i_{t-1} = i]\}, \quad \text{t=1,...,T-1} \tag{6}$$

where

$$C_T(i, I_{t-1}) = \min_{Q_T \geq 0} \{f_T(i, I_{t-1}, Q_T) | i_{T-1} = i]\} \tag{7}$$

represents the boundary condition.

Example. We illustrate the SDP introduced on the same 4-period example in Section 2. We assume $K = 100$, $h = 1$, $b = 10$, and $c = 1$. We observe that the minimised expected total cost is 262.60 when the opening inventory level is 70.

4 MILP-based models

The (R, S) policy features two control parameters: review periods (R), and order-up-to-levels (S). Under this policy, both R and S are determined at the beginning of the planning horizon; an order is issued to reach the order-up-to-level at the beginning of each review period.

In the literature, [11] built an MILP model upon the piecewise linearisation approach for the first order loss function $L(x, \omega)$ and its complementary function $\hat{L}(x, \omega)$, where ω represents an independent random variable with the probability density function $g_\omega(\cdot)$ and x denotes a scalar variable. Consider a partition of the support Ω of ω into W disjoint compact subregions $\Omega_1, \ldots, \Omega_W$. By fixing a priory the probability mass $p_i = Pr\{\omega \in \Omega_i\}$, the associated conditional expectation $\mathrm{E}[\omega|\Omega_i]$ are determined. Based on Jensen's and Edmundson-Madanski inequalities, the first order loss function and its complementary function are approximated with piecewise linear functions $(\sum_{i=1}^{W} p_i L(x, E[\omega|\Omega_i]), \sum_{i=1}^{W} p_i \hat{L}(x, E[\omega|\Omega_i]))$. For a special case of standard normally distributed random variables, all p_i and $E[\omega|\Omega_i]$ are precomputed in [13].

We now consider a correlated random variable d_t, for $t = \{1, \ldots, T\}$, we can compute the conditional distribution $f_{d_t|i_{t-1}}(\cdot)$ of $d_t|i_{t-1}$ based on Lemma 1. We apply the piecewise linear approximation proposed in [13] on its conditional distribution. Therefore, $L(x, d_t|i_{t-1})$ and $\hat{L}(x, d_t|i_{t-1})$ are approximated by $\sum_{i=1}^{W} p_i L\big(x, E\big[\{d_t|i_{t-1}\}|\Omega_i\big]\big)$ and $\sum_{i=1}^{W} p_i \hat{L}\big(x, E\big[\{d_t|i_{t-1}\}|\Omega_i\big]\big)$, respectively.

Example. We illustrate the MILP model introduced on the same example in Section 3. We observe that the minimum expected total cost is 256.07, when the opening inventory level is 70. Specifically, the reviewing time periods are 1 and 3, and the corresponding order-up-to-levels are 69.18 and 114.34.

5 Conclusion

In this paper we presented a MILP-based model for approximating optimal (R, S) policy parameters with correlated demand. This model is based on a mathematical programming model that can be solved by using-off-the-shelf optimization packages. Our preliminary results show that the optimality gap of our model is tighter, and the computational time of our model is reasonable. This model also can be extended to cover time-series-based demand process, such as AR, MA, ARMA, ARCH.

References

[1] Ronald G Askin. A procedure for production lot sizing with probabilistic dynamic demand. *Aiie Transactions*, 13(2):132–137, 1981.

[2] Richard Bellman. Dynamic programming. *Princeton University Press*, 89:92, 1957.

[3] James H Bookbinder and Jin-Yan Tan. Strategies for the probabilistic lot-sizing problem with service-level constraints. *Management Science*, 34(9):1096–1108, 1988.

[4] Fangruo Chen and Jing-Sheng Song. Optimal policies for multiechelon inventory problems with markov-modulated demand. *Operations Research*, 49(2):226–234, 2001.

[5] Feng Cheng and Suresh Sethi. Optimality of state-dependent (s, S) policies in inventory models with markov-modulated demand and lost sales. 1999.

[6] Awi Federgruen and Paul Zipkin. An efficient algorithm for computing optimal (s, S) policies. *Operations research*, 32(6):1268–1285, 1984.

[7] Stergios Fotopoulos, Min-Chiang Wang, and S Subba Rao. Safety stock determination with correlated demands and arbitrary lead times. *European Journal of Operational Research*, 35(2):172–181, 1988.

[8] Stephen C Graves. A single-item inventory model for a nonstationary demand process. *Manufacturing & Service Operations Management*, 1(1):50–61, 1999.

[9] Gordon Duane Johnson and HE Thompson. Optimality of myopic inventory policies for certain dependent demand processes. *Management Science*, 21(11):1303–1307, 1975.

[10] Grzegorz Reyman. State reduction in a dependent demand inventory model given by a time series. *European journal of operational research*, 41(2):174–180, 1989.

[11] Roberto Rossi, Onur A Kilic, and S Armagan Tarim. Piecewise linear approximations for the static–dynamic uncertainty strategy in stochastic lot-sizing. *Omega*, 50:126–140, 2015.

[12] Roberto Rossi, S Armagan Tarim, Brahim Hnich, and Steven Prestwich. A global chance-constraint for stochastic inventory systems under service level constraints. *Constraints*, 13(4):490–517, 2008.

[13] Roberto Rossi, S Armagan Tarim, Steven Prestwich, and Brahim Hnich. Piecewise linear lower and upper bounds for the standard normal first order loss function. *Applied Mathematics and Computation*, 231:489–502, 2014.

[14] Herbert Scarf. The optimality of (s, S) policies in the dynamic inventory problem. 1959.

[15] Jing-Sheng Song and Paul Zipkin. Inventory control in a fluctuating demand environment. *Operations Research*, 41(2):351–370, 1993.

[16] S Armagan Tarim, Mustafa K Dogru, Ulaş Özen, and Roberto Rossi. An efficient computational method for a stochastic dynamic lot-sizing problem under service-level constraints. *European Journal of Operational Research*, 215(3):563–571, 2011.

[17] S Armagan Tarim and Brian G Kingsman. The stochastic dynamic production/inventory lot-sizing problem with service-level constraints. *International Journal of Production Economics*, 88(1):105–119, 2004.

[18] S Armagan Tarim and Brian G Kingsman. Modelling and computing (R_n, S_n) policies for inventory systems with non-stationary stochastic demand. *European Journal of Operational Research*, 174(1):581–599, 2006.

[19] S Armagan Tarim and Barbara M Smith. Constraint programming for computing non-stationary (R, S) inventory policies. *European Journal of Operational Research*, 189(3):1004–1021, 2008.

[20] Arthur F Veinott Jr. Optimal policy for a multi-product, dynamic, nonstationary inventory problem. *Management Science*, 12(3):206–222, 1965.