



THE UNIVERSITY *of* EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e.g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

This work is protected by copyright and other intellectual property rights, which are retained by the thesis author, unless otherwise stated.

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author.

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given.

Mathematical Programming Heuristics for Nonstationary Stochastic Inventory Control

Mengyuan Xiang



A thesis submitted in fulfilment of the requirements for the
degree of Doctor of Philosophy

The University of Edinburgh
2019

Declaration

This thesis has been composed by myself and contains no material that has been accepted for the award of any other degree at any university.

A part of this thesis has been published in the *European Journal of Operational Research*:

Xiang, M., Rossi, R., Martin-Barragan, B., and Tarim, S. A. (2018). Computing non-stationary (s, S) policies using mixed integer linear programming. *European Journal of Operational Research*, 271(2), 490-500.

Permission to include text from that paper has been gained from the publisher and the authors.

To the best of my knowledge and belief this thesis contains no other material previously published by any other person except where due acknowledgement has been made.

(Mengyuan Xiang)



Acknowledgment

First and foremost, I would like to express my deepest gratitude to my supervisor Prof. Roberto Rossi. Many thanks for him for offering me the opportunity to be his PhD student in the summer of 2015. He is the funniest advisor and one of the smartest people I know. I thank him for his generosity in sharing his experiences and observations with me, and always encouraging and believing in me. I sincerely appreciate the opportunity to work with him. Many thanks to Roberto, especially when Belen was on Maternity leave.

I would like to thank my co-supervisor Dr Belen Martin-Barragan, for all her benign guidance and invaluable support during this research, and for having faith in me to be able to achieve this accomplishment. I will never forget those afternoons when we wrote papers together in her office. I am forever gracious for all the time and energy she spent in helping me.

I thank Prof. S. Armagan Tarim. His attitude towards research, enthusiasm and generosity thought me a lot. I will always remember our meeting at the restaurant of Sheraton Poznan Hotel in July 2016. Having a chance to collaborate with him was a privilege for me.

My time in Edinburgh was enjoyable in a large degree due to the presence of my colleagues and friends at the University of Edinburgh Business School, with whom I have shared moments of deep anxiety but also big excitement. I would like to thank all of them for their support and assistance in my four years PhD journey.

My appreciation goes as well as to Prof. Thomas Archibald as internal examiner, and Prof. Eligius M T Hendrix, as external examiner, for their time and efforts in providing insightful and constructive comments and suggestions.

Last but not least, I would like to take this opportunity to thank my parents and grandparents who continuously support me with unique patience during the last four years. I am very lucky and grateful for having such a wonderful family.

Abstract

This work focuses on the computation of near-optimal inventory policies for a wide range of problems in the field of nonstationary stochastic inventory control. These problems are modelled and solved by leveraging novel mathematical programming models built upon the application of stochastic programming bounding techniques: Jensen's lower bound and Edmundson-Madanski upper bound.

The single-item single-stock location inventory problem under the classical assumption of independent demand is a long-standing problem in the literature of stochastic inventory control. The first contribution hereby presented is the development of the first mathematical programming based model for computing near-optimal inventory policy parameters for this problem; the model is then paired with a binary search procedure to tackle large-scale problems.

The second contribution is to relax the independence assumption and investigate the case in which demand in different periods is correlated. More specifically, this work introduces the first stochastic programming model that captures Bookbinder and Tan's static-dynamic uncertainty control policy under nonstationary correlated demand; in addition, it discusses a mathematical programming heuristic that computes near-optimal policy parameters under normally distributed demand featuring correlation, as well as under a collection of time-series-based demand process.

Finally, the third contribution is to consider a multi-item stochastic inventory system subject to joint replenishment costs. This work presents the first mathematical programming heuristic for determining near-optimal inventory policy parameters for this system. This model comes with the advantage of tackling nonstationary demand, a variant which has not been previously explored in the literature.

Unlike other existing approaches in the literature, these mathematical programming models can be easily implemented and solved by using off-the-shelf mathematical programming packages, such as IBM ILOG optimisation studio and XPRESS Optimizer; and do not require tedious computer coding.

Extensive computational studies demonstrate that these new models are competitive in terms of cost performance: in the case of independent demand, they provide the best optimality gap in the literature; in the case of correlated demand, they yield tight optimality gap; in the case of nonstationary joint replenishment problem, they are competitive with state-of-the-art approaches in the literature and come with the advantage of being able to tackle nonstationary problems.

Lay Summary

The aim of stochastic inventory control is to order the right quantity of items, at the right time, in the right location, with the intention of satisfying uncertain demand. This work focuses on tackling unexplored settings in the field of stochastic inventory control by using mathematical programming approaches.

A long-standing problem in stochastic inventory control is to find the replenishment plan for the single-item single-stock location inventory system under the assumption that demand is uncertain over time. The first contribution of this work is to present the first mathematical programming model for solving this class of problems to near-optimality.

Since environmental factors, such as economic conditions, market conditions, and any exogenous conditions, have significant influence on the demand of a product, the second contribution of this work is to investigate the case in which item demand is correlated over time, and to present the first mathematical programming model for tackling the stochastic inventory control problem under correlated demand.

Finally, the third contribution is to consider a multi-item stochastic inventory system in which several products are ordered from the same supplier and incur joint replenishment costs. By ordering products jointly, the ordering cost may be largely reduced. A mathematical programming model is hence presented for tackling the stochastic joint replenishment problem.

Extensive computational studies demonstrate that these new models are competitive in terms of cost performance: in the case of independent demand, they provide the best optimality gap in the literature; in the case of correlated demand, they yield the tight optimality gap; in the case of the joint replenishment problem, they are competitive with state-of-the-art approaches in the literature and come with the advantage of being able to tackle more complex settings.

Contents

Acknowledgement	iii
Abstract	iv
Lay Summary	vi
Contents	vii
List of Figures	x
List of Tables	xii
1 Introduction	1
1.1 Preliminaries	1
1.1.1 Motivations	1
1.1.2 Structure	5
1.2 Formal background	6
1.2.1 Inventory control	6
1.2.1.1 Single-item deterministic inventory control	6
1.2.1.2 Single-item stochastic inventory control	13
1.2.1.3 Multi-item deterministic inventory control	21
1.2.1.4 Multi-item stochastic inventory control	23
1.2.2 Stochastic dynamic programming	26
1.2.3 Stochastic programming bounding techniques	27
1.2.4 Piecewise linear approximation technique	28
1.2.5 Mixed integer linear programming	31
1.3 Related works	35
1.3.1 (s, S) policy	35
1.3.1.1 Single-item inventory system	35
1.3.1.2 Multi-item inventory system	42
1.3.1.3 Multi-echelon inventory system	43

1.3.2	(R, S) policy	43
1.3.3	Joint replenishment problem	48
1.4	Thesis statement	50
1.4.1	Summary	50
1.4.2	Contributions	51
1.4.3	Chapter 2 (Paper I): Computing non-stationary (s, S) policies using mixed integer linear programming	52
1.4.4	Chapter 3 (Paper II): (R, S) policy with correlated demand	53
1.4.5	Chapter 4 (Paper III): Nonstationary (R, S) policies for joint replenishment inventory systems	55
1.5	Future work	58
1.6	Conclusions	60
2	Paper I: Computing non-stationary (s, S) policies using mixed integer linear programming	62
2.1	Introduction	62
2.2	Problem description	65
2.3	The optimality of (s, S) policies in stochastic lot sizing	66
2.4	MINLP approximation of $G_t(y)$ function	67
2.5	An MINLP-based model to approximate (s, S) policy parameters	69
2.6	A binary search approach to approximate (s, S) policy parameters	72
2.7	Computational experiments	74
2.7.1	An 8-period test bed	74
2.7.2	A 25-period test bed	79
2.8	Conclusion	80
	Appendices	81
2.A	The piecewise OPL constraint	81
2.B	The MILP model	83
2.C	Test bed	84
2.D	Model accuracy	85
3	Paper II: (R, S) policy with correlated demand	88
3.1	Introduction	88
3.2	Literature review	90
3.3	A stochastic dynamic programming formulation	93
3.4	Towards an (R, S) policy	95
3.4.1	“Static-dynamic uncertainty” policy	96

3.4.2	Modelling the (R, S) policy via stochastic dynamic programming	97
3.4.3	A stochastic programming model under (R, S) policy	98
3.5	Towards a Mixed Integer Linear Programming model	100
3.5.1	Fixed replenishment cycle problem	100
3.5.2	An MILP model for computing (R, S) policies	101
3.6	Applications to time-series-based demand processes	104
3.7	Computational experiments	105
3.7.1	Multivariate normal distribution	105
3.7.2	Time-series-based demand processes	109
3.8	Conclusion	114
	Appendices	114
3.A	List of symbols	114
3.B	Time series processes	116
3.B.1	AR process	116
3.B.2	MA process	117
3.B.3	ARMA process	118
3.B.4	AR-ARCH process	118
3.C	Test bed	120
4	Paper III: Nonstationary (R, S) policies for joint replenishment inventory systems	121
4.1	Introduction	121
4.2	Literature review	123
4.3	Problem description	126
4.4	An MILP model for approximating non-stationary stochastic (R, S) policies	128
4.4.1	A stochastic program	129
4.4.2	An MILP model	130
4.5	MILP model for approximating the optimal (σ, \vec{S}) policies	136
4.6	Computational Experiments	137
4.7	Conclusion	141
	Appendices	142
4.A	Shortest path reformulation for approximating stationary stochastic (R, S) policies	142
4.B	Computational results	144
	Bibliography	148

List of Figures

1.1	Classification of inventory control	7
1.2	EOQ inventory curve	9
1.3	Inventory curve of EOQ with lead time	10
1.4	The inventory level (\tilde{I}) and inventory position (\tilde{IP}) curve under a (s, Q) policy	14
1.5	The inventory position (\tilde{IP}_t) and inventory level (\tilde{I}_t) curves under a base-stock policy	14
1.6	The expected inventory position curve under (s, S) policy	17
1.7	Plot of K -convexity	19
1.8	The inventory position curve under the (R, Q) policy	19
1.9	The inventory position curve under the (R, S) policy	20
1.10	Two-item EOQ inventory curve	22
1.11	Plot of (σ, \vec{S}) policy	25
1.12	The Jensen's lower bound and the Edmundson-Madansky upper bound in a stochastic minimization problem.	28
1.13	Two-segment piecewise linear lower bound $\mathcal{L}_{lb}(x, Z)$ and upper bound $\mathcal{L}_{ub}(x, Z)$ for the loss function $\mathcal{L}(x, Z)$	32
1.14	Two-segment piecewise linear lower bound $\hat{\mathcal{L}}_{lb}(x, Z)$ and upper bound $\hat{\mathcal{L}}_{ub}(x, Z)$ for the complementary of first order loss function $\hat{\mathcal{L}}(x, Z)$	32
1.15	The syntax of the <code>piecewise</code> command	34
2.1	Plot of $G_1(y)$	67
2.2	The formulation of the non-stationary stochastic lot-sizing problem	68
2.3	Plot of $G_1^s(y)$	70
2.4	Demand patterns in our computational analysis	75
2.5	Expressions for generating demand data	79
2.6	The syntax of the <code>piecewise</code> command in OPL	82
2.7	Rewriting holding and penalty costs by adopting <code>piecewise</code> syntax	83

3.1	A stochastic program under (R, S) policy with correlated demand	99
3.2	An MILP model for computing (R, S) policies with correlated demand	103
3.3	Time series processes	105
3.4	Demand patterns in our computational analysis	107
3.5	Expressions for time-series-based demand patterns	111
4.1	Expected total cost, i.e. $G_1(\vec{I}_0)$, contour plot for the two-item joint replenishment numerical example	129
4.2	Stochastic programming formulation of the JRP.	130
4.3	MILP model for approximating (R, S) policies	134
4.4	Replenishment plans of the 5-item 10-period example	135
4.5	Plot of expected total costs for the two-item joint replenishment numerical example	138
4.6	Expected inventory curve under (R, S) policy.	143
4.7	Shortest path formulation for approximating stationary stochastic (R, S) policies	144

List of Tables

1.1	A summary of literature surveyed on stochastic (s, S) policy . . .	36
1.2	A summary of literature surveyed on single-item stochastic (R, S) policy	44
2.1	Optimal (s, S) policy parameters obtained via the joint MINLP heuristic and the stochastic dynamic programming	72
2.2	Near-optimal (s, S) policy parameters obtained via the binary search approach	74
2.3	Average optimality gaps % of the 8-period numerical experiment for different pivoting parameters	77
2.4	Average computational time(seconds) of the 8-period numerical for different pivoting parameters	78
2.5	BS heuristic on a 25-period test bed, average computational time(seconds) with different parameter settings	80
2.6	Demand data of the 8-period computational analysis	85
2.7	Demand data of the 25-period computational analysis	85
2.8	Model accuracy of the 8-period numerical experiments	86
2.9	Accuracy of the 25-period numerical experiments	87
3.1	Comparison of the MILP model and the stochastic dynamic program	104
3.2	Computational behaviours of the multivariate normally distributed demand processes for different pivoting parameters	108
3.3	Computational times of the multivariate normally distributed demand processes for different pivoting parameters	110
3.4	Computational efficiency of the time-series-based demand processes for different pivoting parameters	112
3.5	Model accuracy of the time-series-based demand processes for different pivoting parameters	113
3.6	A list of symbols	116
3.7	Expected demand data of the 8-period computational analysis . .	120

4.1	Demand rates λ_t^n of the 5-item 10-period example	136
4.2	K^n , λ^n , and L^n of data set Atkins and Iyogun (1988)	139
4.3	Computational results on the data set of Atkins and Iyogun (1988)	139
4.4	Computational results on the data set of Viswanathan (1997) . . .	141
4.5	Computational results on the data set of Atkins and Iyogun (1988)	145
4.6	Computational results on the data set of Viswanathan (1997) . . .	147

Chapter 1

Introduction

1.1 Preliminaries

This section first presents motivations for the work presented in this dissertation, then briefly introduces topics discussed in this dissertation, and lastly outlines the structure of the rest of this chapter.

1.1.1 Motivations

Inventory control aims at ordering the right quantity of items, at the right time, in the right location, with the intention of satisfying demand. Optimising these decisions is becoming increasingly important in nowadays increased globalisation, improved information technology, rapid-updated manufacturing technologies, radical-changed customer behaviour, shortened product life cycles, and quicksilver markets (Elms and Low, 2013).

Business keeps \$1.43 of inventory on hand for every \$1 of sales, and inventory along with accounts receivable and accounts payable has tied up \$1.1 trillion in cash which is equivalent to 7% of the U.S. GDP (Source: REL, 2017). Despite the importance of inventory control, only 33% of small business adopt inventory control systems, nearly 7% never track their inventory at all, 15% use pen and paper, and another 24% use a spreadsheet to track their inventory (Source: WASP Barcode, 2016).

Additionally, recent experiments conducted by GS1 and Auburn University's RFID Lab showed that the average inventory accuracy threshold for retail operations is only 63%. In such a case, items will be out of stock more frequently, customer satisfaction will plummet, and a large amount of capital's liquidity will be blocked.

Fortunately, inventory control has been getting increasing attention. A recent study conducted by the Motorola Future of Warehousing shows that investments in inventory operations technology are increasing; it is predicted that 66% of retailers will have made a significant investment in warehousing and inventory management technology through 2018 (Source: ITE, 2016).

In the past century, the Operations Research community has paid significant attention and developed a large amount of lore to effective inventory control. The first study in this area dates back to Harris (1913), which proposed the well-known *Economic Order Quantity (EOQ)* model to answer the two fundamental research questions

- how large should an order be?
- when should an order be placed?

Since this pioneering work, research on inventory control has been expanding by considering different environments, operating parameters, and modelling assumptions. The solid theoretical foundations upon which the field of inventory control make it one of the most well-developed fields of Operational Research.

One of the fundamental issues in the theory and practice of inventory control has been the modelling of demand uncertainty resulting from inherent qualities of the business and its customer base, or external factors, such as seasonal fluctuations, and customer preference shift. On the one hand, companies will experience leftovers if demand suddenly drops; this will consume physical space, block capital liquidity, and increase the probability of stocks being damaged and lost. On the other hand, companies will not be able to satisfy customer needs if demand suddenly increases; this will increase the probability of poor customer service. Therefore, it is necessary to incorporate demand uncertainty in inventory control.

A pioneering study in stochastic inventory control is Scarf (1960), which first characterised the structure of the optimal control policy for an important class of inventory problems: the single-item single-stock location problem under fixed as well as proportional ordering cost, linear holding and penalty costs. Later, Bookbinder and Tan (1988) proposed new policies for approximating the optimal inventory control policy in Scarf (1960). (Kilic and Tarim, 2011; Tunc et al., 2011; Dural-Selcuk et al., 2016) conducted comparison studies on inventory policies. Thorough literature reviews on the determination of near-optimal policy parameters were conducted by (Aggarwal, 1974; Yano and Lee, 1995; Ullah and Parveen, 2010; Glock et al., 2014; Bushuev et al., 2015).

Existing approaches in the field of stochastic inventory control present several drawbacks.

- *Loose optimality gaps.* Previous research has shown that stochastic inventory control started about sixty years ago, but the computation of optimal policy parameters has not progressed substantially over the past thirty years. As discussed in Dural-Selcuk et al. (2016), this literature still presents loose optimality gaps.
- *Complexity of implementation.* Due to the combinatorial nature of stochastic inventory control, several search-based methods (Bollapragada and Morton, 1999; Özkaya et al., 2006) were presented in the literature to find near-optimal policies; these approaches require dedicated code. Although some easy-to-implement methods have been suggested by (Rossi et al., 2015; Tunc et al., 2018), most existing studies still require considerable implementation efforts.
- *Lack of widely adaptable methodologies.* Most existing methods for stochastic inventory control are ad-hoc, e.g., (Askin, 1981; Bollapragada and Morton, 1999). Due to loose optimality gaps and complexity of implementation of existing approaches, researchers are motivated to find other techniques beyond traditional ad-hoc models. Mathematical programming has been applied in this field, but relevant studies are very limited.

The goal of this work is therefore to develop near-optimal, easy-to-implement, and widely adaptable mathematical programming models for tackling stochastic inventory problems.

Mathematical programming is a widely adopted technique to model complex decision/optimisation problems. Initially introduced during World War II to optimise military operations, it was then transferred to many industries, such as production planning, airline scheduling, resource allocation, and stock and portfolio selection. These industries constantly benefit from the application of mathematical programming. An important subclass of mathematical programming is mixed-integer linear programming (MILP). MILP has been successfully applied in a variety of business areas due to the development of readily available modelling languages and MILP solvers such as GUROBI, IBM ILOG optimisation studio, and XPRESS Optimizer.

Stochastic programming is a technique for modelling optimisation problems that involve uncertainty. The goal of stochastic programming is to find a policy that (i) is feasible for all (or almost all) the possible realisations of uncertain

parameters of a model, and that (ii) optimises the expectation of some function of the decisions and the random variables. Traditional methods to compute lower and upper bounds on the optimal objective value of a stochastic program are to approximate it by a deterministic problem by leveraging Jensen's and Edmundson-Madansky inequalities.

Thesis: mathematical programming based heuristics leveraging Jensen's and Edmundson-Madansky inequalities can be used to tackle a wide range of problems that relax long-standing assumptions in stochastic inventory control; the resulting models are easier to implement than state-of-the-art approaches.

The results discussed in this work fully support this thesis; these approaches here presented are:

- *novel in terms of modelling method.* This work presents new mathematical programming models that are built upon the application of stochastic programming bounding techniques, including Jensen's and Edmundson-Madanski inequalities; it applies these models to solve challenging nonstationary inventory control problems.
- *near optimal.* Results hereby discussed include both state-of-the-art approaches featuring the best optimality gaps in the literature, as well as approaches that - although they do not fully dominate other existing strategies - are competitive in terms of cost performance.
- *easily implemented using standard off-the-shelf mathematical programming packages.* Unlike other existing approaches in the literature, these methods can be implemented and solved by using off-the-shelf mathematical programming packages, such as IBM ILOG optimisation studio and XPRESS Optimizer; and do not require tedious computer coding.
- *broadly applicable.* These modelling approaches relax restrictive assumptions in the stochastic inventory control literature thus addressing unexplored settings, such as computing near-optimal policy parameters under nonstationary demand, tackling correlated demand as well as a collection of time-series-based demand processes, and solving joint replenishment problem under nonstationary demand.

As discussed above, optimal control of complex inventory systems is becoming increasingly important. It is necessary to develop novel, near-optimal, and easily implemented approaches for tackling unexplored settings in the field of stochastic inventory optimisation. The work presented in this dissertation tries to pursue

these objectives.

Topic: *This thesis investigates the application of new mathematical programming based approximations leveraging stochastic programming bounding techniques to a wide range of stochastic inventory control problems. Specifically, it first focuses on the determination of optimal policy parameters of single-item single-stocking location nonstationary stochastic inventory problems under Scarf's setting. It then relaxes the classical assumption of independence of demand distributions and it investigates the case in which demand is correlated between periods; the resulting analysis is extended to a collection of time-series-based demand processes. It finally addresses the case of a multi-item nonstationary inventory system subject to joint replenishment costs. All the resulting models are easily modelled and solved using standard off-the-shelf mathematical programming packages. Extensive computational studies show that these approaches can model and solve optimisation problems that could not be solved or could not be solved exactly by other existing approaches.*

1.1.2 Structure

The rest of this chapter is structured as follows:

- Section 1.2 first presents background information on inventory control, then it discusses stochastic dynamic programming, stochastic programming bounding techniques, piecewise linear approximation techniques, and mathematical programming related topics.
- Section 1.3 conducts a comprehensive literature study on stochastic inventory control, and in particular on stochastic inventory control policies; in addition, this section outlines research gaps in the literature.
- Section 1.4 summarises the content of this thesis, highlights contributions of this work, and presents respective contributions for each of the following chapters.
- Section 1.5 discusses future research directions. Specifically, it first discusses which questions remain open for each of the following chapters, and then presents possible research areas where modelling methods presented in this work can be successfully applied.
- Section 1.6 draws conclusions.

1.2 Formal background

This section first discusses topics in inventory control (Section 1.2.1), then it provides the relevant formal background in stochastic dynamic programming (Section 1.2.2), stochastic programming bounding techniques (Section 1.2.3), and piecewise linear approximation techniques (Section 1.2.4). Finally, mixed integer linear programming related topics are discussed (Section 1.2.5).

1.2.1 Inventory control

This section aims at discussing related topics in inventory control based on the classification in Fig. 1.1. It first discusses the single-item inventory control with deterministic and stochastic demand in Section 1.2.1.1 and Section 1.2.1.2, and then briefly introduces topics in multi-item inventory control including deterministic and stochastic demand in Section 1.2.1.3 and Section 1.2.1.4.

This section is mainly based on (Silver et al., 1998; Zipkin, 2000; Snyder and Shen, 2011).

1.2.1.1 Single-item deterministic inventory control

Inventory control is a very popular research area. It is categorised into *deterministic*, where the demand is known, and *stochastic*, where the demand is unknown but follows a certain type of distribution, based on the nature of demand. It can be further classified into *continuous* and *periodic* based on the inventory review. In continuous-review systems, the inventory is continuously monitored, and an order is placed whenever a certain condition is met; while in periodic-review systems, the inventory is checked every time period, and an order is placed if the reorder condition is met. This section presents classical models in the literature of single-item deterministic inventory control.

There are several measures used to assess the amount of inventory in the system at any given time. Before detailed discussions, three commonly adopted measures are introduced.

- *On-hand inventory*: inventory amount available at the location.
- *Backorder*: demand that has occurred but has not been satisfied because of lack of on-hand inventory.
- *Inventory level (I)*: equals to the *on-hand inventory*, minus the *backorder*. If the inventory level is positive, then items are available in stock; otherwise,




Figure 1.1: Classification of inventory control

no items on-hand are available and one observes backorders.

EOQ model. The oldest and best-known single-item continuous-review deterministic inventory control model is the *economic order quantity* (EOQ) model, proposed by Harris (1913). It is assumed that the demand is deterministic and constant with a rate of λ ; backorders are not allowed; orders are received immediately after being placed. The objective is to determine the optimal order quantity each time an order is placed to minimise the relevant average cost comprising fixed and unit ordering costs, and holding costs.

- *Fixed ordering cost (K):* the cost of placing an order. It is independent of the order quantity and usually accounts for the administrative fee, the delivery cost, and so on.
- *Unit ordering cost (v):* the cost of ordering each unit of the item.
- *Holding cost (h):* the cost of keeping a unit item in inventory. It includes the cost of storage space, taxes, insurance, breakage, opportunity cost, and so on.

Note that, since the total order quantity over the planning horizon is constant in deterministic inventory control, the total unit ordering costs are constant and do not affect the optimal replenishment plan. Therefore, for simplicity, the unit ordering cost is neglected in optimising deterministic inventory control problems.

It is well-known that the optimal solution of the EOQ mode has *zero inventory property* (Theorem 1.2.1), and *constant order sizes* (Theorem 1.2.2), graphically shown in Fig. 1.2.

Theorem 1.2.1 (Zero inventory property). *The optimal replenishment plan of the EOQ model is to place orders when the inventory level is exactly at zero.*

Theorem 1.2.2 (Constant order sizes). *If Q is the optimal order size at time 0, then it will also be the optimal order size every other time an order is placed.*

The optimal solution of the EOQ model is given as follows.

Theorem 1.2.3. *The optimal order quantity Q^* in the EOQ model is given by*

$$Q^* = \sqrt{\frac{2K\lambda}{h}}, \quad (1.1)$$

and the minimised cost per time unit is given by

$$\bar{C}(Q^*) = \sqrt{2Kh\lambda}. \quad (1.2)$$




Figure 1.2: EOQ inventory curve

The optimal order quantity Q^* and the corresponding minimised cost per time unit $\bar{C}(Q^*)$ in Theorem 1.2.3 are achieved under assumptions that orders are received immediately after placing, and backorder is not allowed. However, these assumptions do usually not hold in business practice. Thus, the classical EOQ model has been extended to incorporate fixed lead time and backorders.

EOQ with lead time. In practice, it usually takes hours, days, weeks, or even several months to deliver items from suppliers to customers. The delivery time, known as *lead time* (L) in inventory management literature, is classified into *fixed* or *stochastic*. This work only discusses fixed lead time related topics. Literature on stochastic lead time refers to (Liberatore, 1979; Nasri et al., 1990; Parlar and Berkin, 1991).

In what follows, terminologies are introduced when considering lead times.

- *Outstanding order*: the order that has been placed, but not received by the company at inventory review because of stock inspection, and transportation.
- *Inventory position (IP)*: the sum of *inventory level* and *outstanding order*. When the lead time is neglected, the inventory position is equal to the inventory level.

As shown in Fig. 1.3, the optimal solution of the EOQ model does not change if the fixed lead time is considered; however, decision makers place orders L time units before inventory levels reach 0. It is more convenient to express this term as *reorder point* (s), which accounts for the demand taking place during lead time, i.e. $s = L\lambda$. Therefore, the optimal replenishment plan is to place an order quantity Q^* (Eq. 1.1) when the inventory level reaches reorder point s , and the associated minimised cost per time unit is $\bar{C}(Q^*)$ (Eq. 1.2).




Figure 1.3: Inventory curve of EOQ with lead time

EOQ with backorders. When backorders are allowed in the EOQ model, since demand is constant and deterministic, the number of backorders in every replenishment cycle is the same.

Let x represent the fraction of demand that is backordered. The penalty cost is incurred for every unit unsatisfied demand.

- *Penalty cost (b):* the cost of not having sufficient inventory to meet customer demand. If the excess demand is backordered, the penalty cost includes bookkeeping costs, delay costs, and the loss of goodwill. If excess demand is lost, the penalty cost also includes the lost profit from the missed sale.

The goal is to find the optimal replenishment plan to minimise the total average cost consisting of fixed ordering costs, holding costs, and penalty costs. The optimal solution of EOQ with backorders are given in Theorem 1.2.4.

Theorem 1.2.4. *The optimal solution of the EOQ model with backorders is,*

$$Q^* = \sqrt{\frac{2K\lambda(h+b)}{hb}}, \quad (1.3)$$

$$x^* = \frac{h}{h+b}, \quad (1.4)$$

$$\bar{C}(Q^*, x^*) = \sqrt{\frac{2K\lambda hb}{h+b}}. \quad (1.5)$$

The cost $\bar{C}(Q^*, x^*)$ in Eq. (1.5) is smaller than or equal to $\bar{C}(Q^*)$ in Eq. (1.2), since the classical EOQ is a special case of EOQ with backorders in which $x = 0$. Additionally, the optimal order quantity Q^* in Eq. (1.3) is greater than that in Eq. (1.1), since placing larger orders in the EOQ with backorders does not require to carry quite as much inventory as it does in the classical EOQ. Therefore, the EOQ with backorders offers extra flexibility which allows placing larger orders.

As discussed above, the classical EOQ model and its variants assume constant demand in continuous-review inventory systems. However, the demand usually varies in different time periods in practice, which motivated extensive studies in the determination of optimal ordering plans in time-varying inventory systems. The inventory is checked at the end of each time period, and the replenishment is made at the beginning of a time period if a certain condition is met. The rest of this section surveys two well-known models — *the Wagner-Whitin model* and *Silver-Meal heuristic* — for determining the optimal replenishment plan under time-varying demand.

The Wagner-Whitin model. The Wagner-Whitin model (Wagner and Whitin, 1958) is one of the most famous periodic review inventory models under time-varying demand. Consider a T -period planning horizon, let d_t represent the demand rate of period t , $t = 1, \dots, T$. At the beginning of each time period, the fixed ordering cost K and the unit ordering cost v are incurred if an order is placed. At the end of each time period, the holding cost h is charged for every unit in stock. Backorders are not allowed. Note that the unit ordering cost is neglected since it is a constant.

Additionally, since orders are issued and delivered immediately, similar to the EOQ model, the optimal solution has the *zero inventory property* (Theorem 1.2.1). Therefore, the problem of deciding when and how much to order is equivalent to determine in which periods to order, and the optimal order quantity is the total demand between two consecutive orders.

Let C_t represent the minimised total cost over periods t, \dots, T given an order is placed at the beginning of period t . Then, C_t is defined recursively in terms of C_j for later period j , $j = t, \dots, T + 1$,

$$C_t = \min_{t < j \leq T+1} \left\{ K + h \sum_{i=t}^{j-1} (i-t)d_i + C_j \right\}, \quad (1.6)$$

where the boundary $C_{T+1} = 0$. Note that the first two terms inside the braces account for the total cost incurred over periods $t, \dots, j-1$ given an order is placed in period t ; the last term C_j accounts the cost over periods j, \dots, T when an order is placed in period j .

The Wagner-Whitin algorithm is basically a dynamic programming algorithm (discussed in Section 1.2.2), equivalent to finding the shortest path through the network with $T + 1$ nodes in which each node represent a time period and an arc from period t to period j represents ordering in period t to satisfy demand

of periods $t, t + 1, \dots, j - 1$. The shortest path reformulation has been widely adopted in tackling inventory problems, e.g., (Brahimi et al., 2006; Tarim et al., 2011).

Example. Consider a 4-period example, the demand of each period d_t is 20, 40, 60, and 40, fixed ordering cost $K = 100$, and holding cost $h = 1$. By solving Wagner-Whitin model, the optimal replenishment plan is to order in period 1 and 3, order quantities 60 and 100, respectively, and the minimised total cost is 280.

Silver-Meal heuristic. Although the Wagner-Whitin algorithm provides the optimal replenishment plan, it has some drawbacks from the practitioners point of view, such as the considerable computational effort, the complex nature of the algorithm, and additional assumptions. In this regard, Silver and Meal (1973) proposed a simple variant of the basic EOQ which is commonly adopted in practice.

Let $\bar{C}_{t,j}$ represent the average cost per period over periods $t, \dots, t+j$ assuming that an order is placed at the beginning of period t to cover demand in the next j periods. The optimal replenishment period length j is obtained if $\bar{C}_{t,t+j-1} < \bar{C}_{t,t+j}$. The optimal order quantity Q^* is the demand convolution over periods $t, \dots, t + j - 1$, i.e.: $Q^* = d_t + \dots + d_{t+j-1}$.

Note that this method only guarantees a local minimum of the average cost per time unit for the current replenishment. Since the search procedure is stopped with the first increase in costs per time unit, it is possible to find larger values of j that yield lower costs per time unit. A computational study in Baker (1989) indicated that the Silver-Meal heuristic incurs an average cost penalty for using the heuristic instead of the Wagner-Whitin algorithm of less than 1%. However, computational experiments in Blackburn and Millen (1980) on a rolling horizon setting revealed that the Silver-Meal heuristic outperforms the Wagner-Whitin algorithm.

Example. Consider again the 4-period example of Wagner-Whitin method, the demand of each period d_t is 20, 40, 60, and 40, fixed ordering cost $K = 100$, and holding cost $h = 1$. By adopting the Silver-Meal heuristic, the optimal replenishment plan is to order in period 1 and 3, order quantities 60 and 100, and the associated total cost is 280.

For an overview of literature on deterministic inventory control refers to (Silver, 1981; Pentico and Drake, 2011; Drake and Marley, 2014), recent developments refer to (Cobb, 2016; Dobson et al., 2017; Pervin et al., 2018).

1.2.1.2 Single-item stochastic inventory control

Early works in the literature of inventory control generally assume that the demand is deterministic, as the EOQ model and its variants, the Wagner-Whitin algorithm, and the Silver-Meal heuristic discussed in Section 1.2.1.1. However, in business practice, demand arrives at random, which motivated extensive research in coping demand uncertainty with inventory control, namely stochastic inventory control.

A key concept in stochastic inventory control is a policy, which is basically a rule that provides solutions to an inventory problem. When using policies, a policy should be chosen first and then solving the inventory problem is to calculate the policy parameters. This dissertation discusses the (s, Q) policy, the base-stock policy, and the (s, S) , (R, Q) , and (R, S) policies proposed by Bookbinder and Tan (1988).

Before detailed discussions, there are two important terminologies introduced in stochastic inventory control.

- *Cycle stock*: the amount of on-hand stock which is used to satisfy the expected demand.
- *Safety stock*: the extra inventory on-hand to buffer against demand uncertainty.

(s, Q) policy. This continuous review policy features two control parameters: reorder point s and order quantity Q . Under this policy, decision makers place an order of size Q whenever the inventory position falls below the reorder point s , as shown in Fig. 1.4. It should be noted that the inventory position, rather than the inventory level, is used to trigger an order. This is because the inventory position includes the outstanding order and it takes proper account of the material requested but not yet received.

Base-stock policy. The base-stock policy is also known as the order-up-to policy. It is widely adopted in both continuous and periodic inventory review systems. However, this dissertation only discusses its application in the periodic inventory review system. Under this policy, decision makers observe the current inventory position at the beginning of each time period t and then place an order to bring the inventory position up to S_t , where S_t is a constant and known as the base-stock level (Fig. 1.5).

Newsvendor problem. The newsvendor problem is one of the classic issues in the stochastic inventory optimisation (Arrow et al., 1951; Dvoretzky et al.,




Figure 1.4: The inventory level (\tilde{I}) and inventory position (\tilde{IP}) curve under a (s, Q) policy




Figure 1.5: The inventory position (\tilde{IP}_t) and inventory level (\tilde{I}_t) curves under a base-stock policy

1952). Key insights stemming from an analysis of this problem have wide-ranging implications from managing inventory decisions for organisations in, for example, the airline, and fashion goods industries.

The newsvendor problem is the problem faced by a news vendor who needs to order newspapers in the early morning. A unit ordering cost is charged with buying newspaper from suppliers; revenue is received for selling each piece of newspaper. At the end of the day, a salvage value is received of selling unsold newspapers back to the supplier. If the newsvendor orders too many, some newspapers will have to be scrapped at the end of the day. If the newsvendor does not order enough newspapers, some customer need will not be satisfied and profit will be lost. The goal is to find the optimal number of newspapers to order that

will maximise the expected profit given that the demand is uncertain.

As usual, let h and b represent the holding cost and penalty cost. In the newsvendor problem, the holding cost h typically consists of the purchase cost of each unit, minus any salvage value, but may also include other costs, such as processing costs. Note that since inventory cannot be carried to the next period, this cost is not technically a holding cost, although this work will refer to it this way. Similarly, the penalty cost consists of the selling price, minus the unit ordering cost. Therefore, the goal of the original newsvendor problem is equivalent to finding the optimal order quantity that minimises the total cost, this is equivalent to the more optimistic view of maximising profit.

It is assumed that demand d is a random variable defined by the distribution function $g(d)$ and cumulative distribution function $G(d)$ with estimates of the parameters of the distribution. Let $C(Q)$ represent the expected total cost at the end of the day with order quantity Q at the beginning of the day. Thus, $C(Q)$ can be formulated as,

$$\begin{aligned} C(Q) &= h \cdot \text{E}[\max(Q - d, 0)] + b \cdot \text{E}[\max(d - Q, 0)] \\ &= h \int_0^Q (Q - d)g(d)d(d) + b \int_Q^\infty (d - Q)g(d)d(d). \end{aligned} \quad (1.7)$$

Consider a random variable ω and a scalar variable x . The first order loss function $\mathcal{L}(x, \omega)$ is defined as,

$$\mathcal{L}(x, \omega) = \int_x^{+\infty} (t - x)g_\omega(t)d(t), \quad (1.8)$$

and its complementary function $\hat{\mathcal{L}}(x, \omega)$ is defined as,

$$\hat{\mathcal{L}}(x, \omega) = \int_{-\infty}^x (x - t)g_\omega(t)d(t). \quad (1.9)$$

Therefore, Eq. (1.7) can be written as follows, by means of the first order loss function and its complementary function,

$$C(Q) = h\hat{\mathcal{L}}(Q, d) + b\mathcal{L}(Q, d). \quad (1.10)$$

Consider a continuous distribution for d , the optimal quantity Q^* is obtained by taking the first order derivative of $C(Q)$ and setting it to zero. Note that $C(Q)$ is convex since $\hat{\mathcal{L}}(Q, d)$ and $\mathcal{L}(Q, d)$ are convex, which guarantees $C(Q^*)$ is a global minimum. This also can be done by checking the second order derivative

is non-negative.

Theorem 1.2.5. *The optimal order quantity Q^* of the Newsvendor problem is given by*

$$Q^* = G^{-1}\left(\frac{b}{h+b}\right). \quad (1.11)$$

The ratio $\frac{b}{h+b}$ is known as α *service level*, i.e., the probability to satisfy customer demand without facing any backorders or lost sales. In business practice, decision makers often pre-define a service level and then compute the optimal order quantity accordingly. The optimal order quantity Q^* in the Newsvendor problem can be interpreted as the minimum number of newspapers to satisfy all customers with probability $100\alpha\%$.

This seminal problem in stochastic inventory control provides useful intuition and a useful decision-making tool, especially, in balancing the holding cost and the penalty cost. In what follows, this dissertation considers a special case of the Newsvendor problem where the demand is normally distributed.

Normal demand distribution. Let d be a normally distributed random variable, $d \sim \mathcal{N}(\mu, \sigma)$, with probability density function $\phi(\cdot)$ and cumulative density function $\Phi(\cdot)$. Let z_α denote the α th fractile of the standard normal distribution, i.e., $z_\alpha = \Phi^{-1}(\alpha)$. According to Theorem 1.2.5, the optimal order quantity is,

$$\begin{aligned} Q^* &= \Phi^{-1}\left(\frac{b}{h+b}\right) \\ &= \mu + z_\alpha \sigma. \end{aligned} \quad (1.12)$$

The first term μ in Eq. (1.12) represents the cycle stock which is used to meet the expected demand; the second term $z_\alpha \sigma$ can be interpreted as the safety stock which is to buffer against demand uncertainty. It is clear that the Newsvendor problem is a special case of a base-stock policy where the inventory level in the early morning of that day is zero.

Example. This work illustrates concepts introduced above on the following example. It is assumed that $d \sim \mathcal{N}(100, 20)$, $h = 0.2$, and $b = 0.8$. By applying Eq. (1.12), to obtain the optimal order quantity is 117, and the associated expected total cost is 16.

In what follows, the single-period Newsvendor problem is extended to multi-period where the inventory leftover will be carried to the next time period. The (s, S) , (R, Q) , and (R, S) policies originally introduced in Bookbinder and Tan

(1988) are presented in the rest of this section to effectively tackle this class of stochastic inventory problems.

(s, S) policy. Under this periodic-review policy, the reorder point s_t , and order-up-to position S_t are fixed for each period t at the beginning of the planning horizon. Decision makers order up to S_t if the inventory position at the beginning of period t is less than the reorder point s_t (Fig. 1.6).

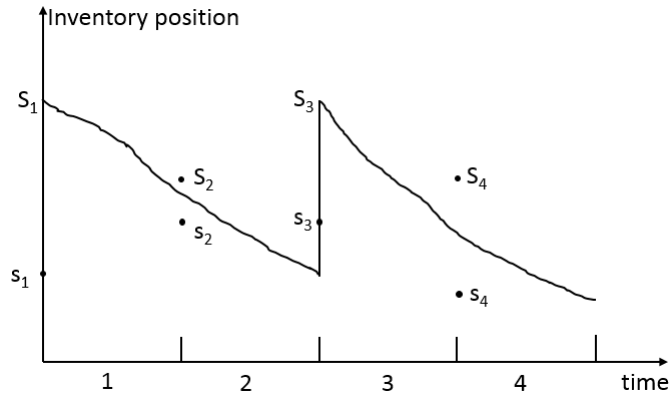


Figure 1.6: The expected inventory position curve under (s, S) policy

Policy optimality. Consider a dynamic problem where the ordering cost consists of fixed and unit ordering costs, and holding and penalty costs are linear. Scarf (1960) proved the optimal policy for this class of problem is always of the (s, S) type.

Let $c(Q_t)$ represent the ordering cost of period t , $t = 1, \dots, T$, with order quantity Q_t ,

$$c(Q_t) = \begin{cases} K + c Q_t, & Q_t > 0 \\ 0, & Q_t = 0. \end{cases}$$

Let $C_t(I_{t-1})$ denote the minimised expected total cost of period t given opening inventory I_{t-1} ,

$$C_t(I_{t-1}) = \min_{Q_t} \left\{ c(Q_t) + h\hat{\mathcal{L}}(I_{t-1} + Q_t, d_t) + b\mathcal{L}(I_{t-1} + Q_t, d_t) + \mathbb{E}[C_{t+1}(I_{t-1} + Q_t - d_t)] \right\}, \quad (1.13)$$

where “E” represents the expectation operator, and the boundary condition

$$C_T(I_{T-1}) = \min_{Q_T} \left\{ c(Q_T) + h\hat{\mathcal{L}}(I_{T-1} + Q_T, d_T) + b\mathcal{L}(I_{T-1} + Q_T, d_T) \right\}. \quad (1.14)$$

Scarf (1960) proved the optimality of the (s, S) policy based on the study of the function

$$G_t(y) = cy + h\hat{\mathcal{L}}(y, d_t) + b\mathcal{L}(y, d_t) + E[C_{t+1}(y - d_t)], \quad (1.15)$$

where y is the inventory level immediately after an order is received.

In order to demonstrate the optimality of the (s, S) policy, the following definition and lemmas are introduced.

Definition 1.2.1 (K-convexity (Scarf, 1960)). Let $K \geq 0$, then $f(x)$ is K -convex if

$$K + f(a + x) - f(x) - a \left\{ \frac{f(x) - f(x - b)}{b} \right\} \geq 0,$$

for all positive a, b and x .

Lemma 1.2.6. *Let f be a continuous, K -convex function. Let S^* be its smallest global minimizer and s^* be the largest $x \leq S^*$ such that $f(x) = f(S^*) + K$.*

Lemma 1.2.7. *The following properties were introduced in Scarf (1960):*

- *0-convexity is equivalent to ordinary convexity.*
- *If $f(x)$ is K -convex, then $f(x + \epsilon)$ is K -convex for all constants ϵ .*
- *If f and g are K -convex and M -convex respectively, then $\alpha f + \beta g$ is $(\alpha K + \beta M)$ -convex when α and β are positive.*

Scarf's proof can be briefly described as follows. According to Lemma 1.2.7, firstly, the boundary $G_T(y) = cy + h\hat{\mathcal{L}}(y, d_T) + b\mathcal{L}(y, d_T)$ is convex since every element on the right-hand-side is K -convex; then, the boundary $C_T(I_{T-1}) = \min_{Q_t} \{-cI_{T-1} + G_T(I_{T-1})\}$ is K -convex; Scarf recursively proved that $G_{T-1}(y)$, $C_{T-1}(I_{T-2}), \dots, G_t(y)$ are K -convex. By applying Lemma 1.2.6, there exists a S^* minimising $G_t(S^*)$, and a unique $s \leq S^*$ such that $G_t(s) = G_t(S^*) + K$. Therefore,

$$C_t(I_{t-1}) = \begin{cases} -cI_{t-1} + G_t(I_{t-1}), & s \leq I_{t-1} \leq S^* \\ -cI_{t-1} + G_t(S^*) + K, & 0 \leq I_{t-1} < s. \end{cases} \quad (1.16)$$

This is sufficient to demonstrate the K -convexity of $C_t(I_{t-1})$ (Fig. 1.7).

Although Scarf (1960) proved the optimality of the (s, S) policy, only two approaches (Askin, 1981; Bollapragada and Morton, 1999) are currently available




Figure 1.7: Plot of K -convexity

in the literature for computing optimal policy parameters. A recent study Dural-Selcuk et al. (2016) shows that these existing studies yield fairly wide optimality gaps in terms of cost performance. To fill this gap in the literature, this work presents models that substantially improve existing optimality gaps; this will be discussed in Chapter 2 (Paper I).

(R, Q) policy. This policy features two control parameters: time interval between two consecutive replenishments R , and quantity of replenishment Q . By operating under this policy, the replenishment period t and the corresponding order quantity Q_t are fixed at the beginning of the planning horizon. At the beginning of each replenishment period t , decision makers place an order with ordering quantity Q_t (Fig. 1.8).




Figure 1.8: The inventory position curve under the (R, Q) policy

Under the (s, S) policy, the inventory position is checked at the beginning of each period, but the actual order quantity is decided when the opening inventory

position is realised. The (R, Q) policy requires decision makers to determine the size and timing of replenishment at the beginning of the planning horizon.

(R, S) policy. Under this policy, the replenishment cycle length R , and order-up-to position S , are determined at the beginning of the planning horizon. Decision makers order up to S_t at the beginning of replenishment period t (Fig. 1.9).




Figure 1.9: The inventory position curve under the (R, S) policy

The static (R, Q) policy is not sufficiently flexible; while the dynamic (s, S) policy suffers from “nervousness” of control action (Kilic and Tarim, 2011; Tunc et al., 2013); the (R, S) policy features characteristics of both the (R, Q) and (s, S) policies, i.e., the timing of replenishment is pre-fixed, while the actual replenishment quantity depends on demand realisation.

The (s, S) policy has been proved to be optimal for a dynamic problem where the ordering cost consists of fixed and unit ordering costs, and holding and penalty costs are linear, provided that we set the parameters of that policy optimally (Scarf, 1960). The (R, S) policy has been showed to have the potential to replace the cost-optimal (s, S) policy for systems with limited flexibility (Kilic and Tarim, 2011). Moreover, it has advantages in organising joint replenishment and shipment consolidation (Tempelmeier, 2013; Silver et al., 1998; Relvas et al., 2013). The (R, Q) policy is appealing in material requirement planning systems, for which order synchronisation is a key concern (Kilic and Tarim, 2011).

Additionally, the demand also can be classified into *independent* and *dependent* in stochastic inventory management literature. Above discussions generally assume that the demand is independent identically distributed with known distribution parameters, such as normal distribution, and Poisson distribution. However, Song and Zipkin (1993) pointed out that 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 regard, many studies on dependent demand processes have emerged, e.g., (Iglehart, 1962; Sethi and Cheng, 1997; Johnson and Thompson, 1975; Dong and Lee, 2003; Carrizosa et al., 2016).

Literature on dependent demand can be further divided into four categories: Markov-modulated, time-series-based, stock-dependent, and time-proportional. This work is limited to model the dependent demand as time-series-based processes, such as Autoregressive (AR), Moving average regressive (MA), autoregressive Moving-Average regressive (ARMA), and Autoregressive with autoregressive conditional heteroskedasticity (AR-ARCH) processes; these processes are further discussed in Chapter 3 (Paper II). Literature on Markov-modulated demand refers to (Fabens and Karlin, 1960; Iglehart, 1962; Song and Zipkin, 1993), stock-dependent demand refers to (Padmanabhan and Vrat, 1995; Hou, 2006; Wu et al., 2006), and time-proportional demand refers to (Silver, 1979; Xu and Wang, 1990; Datta and Pal, 1991).

1.2.1.3 Multi-item deterministic inventory control

This section discusses related topics in multi-item inventory control. It can be divided into *single-echelon* where items are held in stock at the same location, and they share a common supplier or mode of transportation, and *multi-echelon* where items are held in stock at different locations. The inventory control of the single-echelon system focuses on determining the appropriate inventory for each individual unit within the supply chain; while that of the multi-echelon system aims at managing inventories across the entire supply chain, where inventories at each echelon of the supply chain have an impact on the required inventories at different echelons. The focus of this work on the single-echelon multi-item inventory problem, particularly, *joint replenishment problem* (JRP); multi-echelon inventory problems could refer to (Silver et al., 1998; Zipkin, 2000).

Regarding the JRP, every time a group order is placed, the group fixed ordering cost K is incurred regardless the number of items replenished. Moreover, there are also item-specific fixed ordering costs k_n that are charged whenever an item is included in a replenishment order. The goal of the JRP is to determine the optimal inventory replenishment plan minimising the total cost of replenishing multiple items.

EOQ with multiple items. Consider the N -item JRP with constant demand

(Fig. 1.10), all assumptions behind the derivation of the single-item EOQ are retained, except that coordination of items is allowed to reduce ordering costs. Let T denote the group cycle time, and $m_n T$ denote the cycle time of item n , where m_n is an integer multiplier, and $n = 1, \dots, N$. That is item n will be replenished every m_n replenishment of the group. Thus, the order quantity Q_n for item n is $\lambda_n m_n T$. The average inventory level of item n is $\frac{\lambda_n m_n T}{2}$. Therefore, the total cost per unit time is given by

$$C(T, m_n) = \frac{K + \sum_{n=1}^N \frac{k_n}{m_n}}{T} + h \sum_{n=1}^N \frac{\lambda_n m_n T}{2}. \quad (1.17)$$




Figure 1.10: Two-item EOQ inventory curve

Arkin et al. (1989) proved that the JRP is an NP-hard problem even under deterministic demand. Therefore, it is unlikely that there exists a polynomial time algorithm to solve this problem. Algorithms and heuristics for solving deterministic JRP are discussed in (Silver, 1976; Kaspi and Rosenblatt, 1983; Viswanathan, 1996; Fung and Ma, 2001; Viswanathan, 2002).

1.2.1.4 Multi-item stochastic inventory control

In the deterministic joint replenishment inventory system, demand for each individual item is known to be constant. The problem is to determine the length of replenishment cycles, and the optimal order quantity of each item which is the total demand until the next replenishment arrives. In the stochastic joint replenishment inventory system the demand for individual items is unknown but follows a certain type of distribution. The problem is to decide the optimal parameters of a given inventory policy. The focus of Chapter 4 (paper III) will be this problem. This section introduces existing policies in the literature of stochastic JRPs.

(s, c, S) policy. This continuous review policy was introduced by Balintfy (1964). It features three control parameters: reorder point s , can-order level c , and order-up-to position S . Under this policy, When the inventory position of an item i crosses s_i , a replenishment order is triggered to raise its inventory position to S_i ; meanwhile, any other item j with an inventory position at or below its can-order point, $c_j(s_j < c_j < S_j)$, is also included in the replenishment, raising its inventory position to S_j .

(Q, S) policy. This continuous review policy was first proposed by Renberg and Planche (1967). By operating under this policy, whenever the total inventory position drops to the group reorder point, an order is placed to raise the inventory position of each item to item-specific order-up-to position S . The combined order quantity is Q , and the group reorder point is reached when the combined demand reaches Q .

$Q(s, S)$ policy. This continuous review policy is proposed by (Nielsen and Larsen, 2005). By operating under this policy, the total inventory position is continuously evaluated while item-specific inventory positions are reviewed only when the total consumption since the last order reaches Q . Then, every item with inventory position less than or equal to its respective reorder point s is ordered up to level S .

(Q, S, T) policy. This continuous review policy is proposed by Özkaya et al. (2006). Under this policy, decision makers raise the inventory position of each item n to its order-up-to position S_n whenever a total of Q demand accumulated or T time units have elapsed, whichever occurs first.

(σ, \vec{S}) policy. Under this periodic review policy, decision makers order up to \vec{S} if opening inventory levels $\vec{I} \in \sigma$ and $\vec{I} \leq \vec{S}$ ($\vec{S} \in \mathcal{R}^N$, N represents the number

of items); and do not to order, otherwise. The definition of σ is general; its shape and properties are literally unknown. There is no guarantee of σ by convex, or even connected.

Optimality of (σ, \vec{S}) policy. Since the landmark study of Scarf (1960) which proved the optimality of the (s, S) policy for the single-item inventory system, there have been few attempts to prove the optimality for multi-item inventory systems. Gallego and Sethi (2005) gave the general definition of K -convexity in \mathcal{R}^N , and developed properties of K -convex functions which provide solutions to JRPs with the cases of both joint setup and individual setup costs.

Definition 1.2.2. Function $f(\cdot) : \mathcal{R}^N \rightarrow \mathcal{R}$ is K -convex if

$$f(ax + (1 - a)z) \leq af(x) + (1 - a)[f(z) + \mathbf{K}\delta(z - x)],$$

where $x \leq z$, $a \in [0, 1]$, and $\mathbf{K}\delta(z - x)$ is defined as follows,

$$\mathbf{K}\delta(z - x) = K\delta(e'x) + \sum_{n=1}^N k^n \delta(x_n),$$

where $e' = (1, 1, \dots, 1)' \in \mathcal{R}^N$, $\delta(0) = 0$, and $\delta(y) = 1$ for all $y > 0$.

Gallego and Sethi (2005) derived the optimal policy for the joint setup cost case by studying function

$$G_t(\vec{y}) = h\hat{\mathcal{L}}(\vec{y}, \vec{d}_t) + b\mathcal{L}(\vec{y}, \vec{d}_t) + \mathbb{E}[C_{t+1}(\vec{y} - \vec{d}_t)], \quad (1.18)$$

where vector $\vec{y} = (y^1, \dots, y^N)$, and $\vec{d}_t = (d_t^1, \dots, d_t^N)$.

Consider a continuous K -convex function $G_t(\cdot)$, then it has global minimum at \vec{S}_t . Define set $\Sigma = \{\vec{I}_{t-1} \leq \vec{S}_t | G_t(\vec{I}_{t-1}) \leq G_t(\vec{S}_t) + K\}$, and set $\sigma = \{\vec{I}_{t-1} \leq \vec{S}_t | \vec{I}_{t-1} \notin \Sigma\}$. Then, the optimal replenish plan is to order up to \vec{S}_t if opening inventory levels $\vec{I}_{t-1} \in \sigma$ and $\vec{I}_{t-1} \leq \vec{S}_t$; and not to order, otherwise (Gallego and Sethi, 2005).

Due to the complexity of the (σ, \vec{S}) policy, the computation of optimal policy parameters has not been developed. This work shows that new mathematical programming based models can be used to determine whether a given initial inventory level \vec{I}_0 belongs to σ . Details refer to Chapter 4 (Paper III).

Example. The following two-item example illustrates concepts discussed above. Consider an instance in which the group fixed ordering cost is $K = 10$, the item-specific ordering cost k is 0, the holding cost is $h = 1$, the stock-out

penalty cost is $p = 5$. The inventory is controlled for two items over a planning horizon of $T = 4$ periods. Let d_t^n denote the random demand for item n in period t , which follows a Poisson distribution with rate λ_t^n ; where $\lambda_t^1 = \lambda_t^2 = 3, 6, 9, 6$.

Assuming the initial inventory level $\vec{I}_0^1 \in \{0, \dots, 6\}$, and $\vec{I}_0^2 \in \{0, \dots, 6\}$. The set σ and Σ are plotted in Fig. 1.11. The optimal policy is to place an order up to $\vec{S} = (5, 5)$ whenever the inventory level vector $\vec{I}_0 = (I_0^1, I_0^2)$ falls in set σ , and not to place an order if \vec{I}_0 falls in set Σ .




Figure 1.11: Plot of (σ, \vec{S}) policy

(R, T) policy. Atkins and Iyogun (1988) proposed two periodic-review (R, T) -type policies, namely periodic policy P and modified periodic policy MP , which differ only in the way ordering periods T_n are determined. Under this policy, every T_n periods, the inventory position of item n is raised to R_n . Note that this policy is equivalent to previous introduced (R, S) policy when the demand is stationary. Details refer to Chapter 4 (Paper III).

$P(s, S)$ policy. Viswanathan (1997) proposed the periodic-review $P(s, S)$ policy, in which the inventory position of each item is reviewed at every fixed and constant time interval. At each review time, the (s, S) policy is applied to each item, so that any item with inventory position at or below s is ordered up to S .

For a thorough review of literature on the joint replenishment problem refer to (Silver and Peterson, 1985; Goyal and Satir, 1989; Van Eijs et al., 1992; Khouja and Goyal, 2008; Bastos et al., 2017). Even though various inventory control policies and methods for computing their parameters were proposed in

the literature, they could only tackle stationary demand. Chapter 4 (Paper III) introduces the first model to capture nonstationary demand.

1.2.2 Stochastic dynamic programming

This section presents stochastic dynamic programming (SDP) which is used as a benchmark for tackling stochastic inventory problems in this dissertation. The discussion is mainly based on Bellman (1966).

The aim of SDP is to determine a policy that minimises the expected total cost incurred over a given planning horizon for a stochastic optimisation problem.

Consider a discrete system defined on T stages in which each stage $t = 1, \dots, T$ is characterised by,

- *state*, $x_t \in X_t$, where X_t is a finite set of feasible states at the beginning of stage t ;
- *action*, $a_t \in A_t$, where A_t is a finite set of feasible actions at the beginning of stage t given state x_t ;
- *expected immediate cost*, $g_t(x_t, a_t)$, represents the cost at the end of stage t given state x_t and action a_t ;
- *transition probability*, $Pr(x_{t+1}|x_t, a_t)$, denotes the probability that leads the system to state x_{t+1} given state x_t and action a_t ;
- *objective function*, $f_t(y_t)$, represents the optimal expected total cost obtained by following an optimal policy over stages t, \dots, T given state y_t . It takes the following form,

$$f_t(y_t) = \min_{a_t \in A_t} \{g_t(y_t, a_t) + \sum_{y_{t+1} \in X_{t+1}} Pr(y_{t+1}|y_t, a_t) f_{t+1}(y_{t+1})\}, \quad (1.19)$$

and

$$f_T(y_T) = \min_{a_T \in A_T} \{g_T(y_T, a_T)\} \quad (1.20)$$

represents the boundary condition of the system.

SDP is a general method aiming at solving stochastic optimisation problems. It is broadly used in operations research, as many of the problems faced in this field deal with decision making under uncertainty. It guarantees an optimal solution. However, it is computationally inefficient due to the so-called ‘‘curse of

dimensionality” (Bellman and Dreyfus, 2015). To overcome this limitation, several approximations of SDP have been discussed in (Si et al., 2004; Powell, 2007, 2009).

1.2.3 Stochastic programming bounding techniques

Stochastic programming is a widely used approach for modelling optimisation problems that involve uncertainty. This section presents stochastic programming and its bounding techniques—Jensen’s lower bound and Edmundson-Madansky upper bound, which can be used to construct piecewise linear approximations.

This section is mainly based on (Birge and Louveaux, 2011; Kall et al., 1994).

Similar to SDP, recall that $g_t(x_t, a_t)$ denotes the expected immediate cost of stage t given state x_t and action a_t . Then, a T -stage stochastic programming can be written in the following general formulation,

$$\min \sum_{t=1}^T g_t(x_t, a_t), \quad (1.21)$$

where state x_t follows the transition function $x_{t+1} = F^\pi(x_t, a_t, W_{t+1})$, and W_{t+1} represents the realised exogenous information at the beginning of $t + 1$.

The traditional method to compute lower and upper bounds on the optimal objective values of a stochastic program is to formulate a deterministic problem by replacing all the random variable by their expected values and to use bounding techniques — *Jensen’s lower bound* and *Edmundson-Madansky upper bound*.

Jensen’s lower bound. Consider a convex function $\phi(\zeta)$ defined on support $\Omega = [a, b]$. This function can be bounded from below by a linear function $L(\zeta)$ (Theorem 1.2.8), as shown in Fig. 1.12. Additionally, the best lower bound is tangent to $\phi(\zeta)$ at point $E[\zeta]$ (Kall et al., 1994).

Theorem 1.2.8. *If $\phi(\zeta)$ is convex, then*

$$E[\phi(\zeta)] \geq \phi(E[\zeta]).$$

Edmundson-Madansky upper bound. Consider a convex function $\phi(\zeta)$ defined on support $\Omega = [a, b]$, and a linear function $U(\zeta)$ between the two points $(a, \phi(a))$ and $(b, \phi(b))$, as shown in Fig. 1.12. This linear function $U(\zeta)$ is an

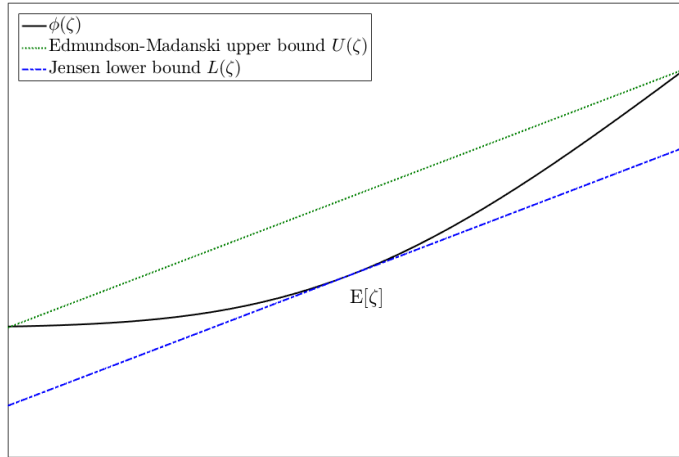


Figure 1.12: The Jensen's lower bound and the Edmundson-Madansky upper bound in a stochastic minimization problem.

upper bound of function $\phi(\zeta)$ and can be represented by

$$U(\zeta) = \frac{\phi(b) - \phi(a)}{b - a}\zeta + \frac{b}{b - a}\phi(a) - \frac{a}{b - a}\phi(b).$$

Jensen's lower bound and Edmundson-Madansky upper bound are common approaches for approximating the objective values of a stochastic program. One way to visualise this lower bound is to assume that the probability distribution in the problem is replaced by a degenerate distribution that put mass only on the expected values of the random variables. The Edmundson-Madansky upper bound is obtained by replacing the probability distribution in the problem with two point distributions that put mass only on the extreme points of the support of the random variables. These bounding techniques have been extended to many dimensions (Frauendorfer, 1996; Kuhn, 2006; Natarajan and Teo, 2017).

1.2.4 Piecewise linear approximation technique

This section applies stochastic programming bounding techniques—Jensen's lower bound and Edmundson-Madansky upper bound, to approximate the first order and its complementary function by piecewise linear functions.

This section is mainly based on Rossi et al. (2014).

Recall the first order loss function

$$\mathcal{L}(x, \omega) = \int_x^{+\infty} (t - x)g_\omega(t)d(t)$$

and its complementary function

$$\hat{\mathcal{L}}(x, \omega) = \int_{-\infty}^x (x-t)g_{\omega}(t)d(t),$$

they have a close relationship.

Lemma 1.2.9. *The first order loss function $\mathcal{L}(x, \omega)$ can be expressed as*

$$\mathcal{L}(x, \omega) = \hat{\mathcal{L}}(x, \omega) - (x - \tilde{\omega}), \quad (1.22)$$

where $\tilde{\omega}$ denotes the expected value of random variable ω .

In what follows, this work presents lower bounds of $\mathcal{L}(x, \omega)$ and $\hat{\mathcal{L}}(x, \omega)$ by using Jensen's inequality (Theorem 1.2.8).

Consider a partition of the support Ω of ω into W disjoint compact subregions $\Omega_1, \dots, \Omega_W$. Let $g_{\omega}(\cdot)$ represent the probability density function of ω . This work defines, for $i = 1, \dots, W$,

$$p_i = Pr\{\omega \in \Omega_i\} = \int_{\Omega_i} g_{\omega}(t)dt,$$

and

$$E[\omega|\Omega_i] = \frac{1}{p_i} \int_{\Omega_i} t g_{\omega}(t)dt.$$

Let $\hat{\mathcal{L}}_{lb}(x, \omega)$ denote the lower bound of the complementary of the first order loss function $\hat{\mathcal{L}}(x, \omega)$, by applying Theorem 1.2.8,

$$\hat{\mathcal{L}}_{lb}(x, \omega) = \sum_{i=1}^W p_i \max(x - E[\omega|\Omega_i], 0). \quad (1.23)$$

This function is equivalent to

$$\hat{\mathcal{L}}_{lb}(x, \omega) = \begin{cases} 0 & -\infty \leq x \leq E[\omega|\Omega_1] \\ p_1 x - p_1 E[\omega|\Omega_1] & E[\omega|\Omega_1] \leq x \leq E[\omega|\Omega_2] \\ (p_1 + p_2)x - (p_1 E[\omega|\Omega_1] + p_2 E[\omega|\Omega_2]) & E[\omega|\Omega_2] \leq x \leq E[\omega|\Omega_3] \\ \dots & \dots \\ (p_1 + p_2 + \dots + p_W)x - (p_1 E[\omega|\Omega_1] + p_2 E[\omega|\Omega_2] + \dots + p_W E[\omega|\Omega_W]) & E[\omega|\Omega_W] \leq x \leq +\infty \end{cases} \quad (1.24)$$

which is piecewise linear in x with breakpoints at $E[\omega|\Omega_1], E[\omega|\Omega_2], \dots, E[\omega|\Omega_W]$.

Lemma 1.2.10. Consider the i -th segment of $\hat{\mathcal{L}}_{lb}(x, \omega)$

$$\hat{\mathcal{L}}_{lb}^i(x, \omega) = x \sum_{k=1}^i p_k - \sum_{k=1}^i p_k E[\omega|\Omega_k] \quad E[\omega|\Omega_i] \leq x \leq E[\omega|\Omega_{i+1}], \quad (1.25)$$

where $i = 1, \dots, W$, and the 0-th segment is $x = 0$, $-\infty \leq x \leq E[\omega|\Omega_1]$.

Lemma 1.2.11. The i -th segment of $\mathcal{L}_{lb}(x, \omega)$ can be written as Eq. (1.26), by applying Lemma 1.2.9,

$$\mathcal{L}_{lb}^i(x, \omega) = x \sum_{k=1}^i p_k - \sum_{k=1}^i p_k E[\omega|\Omega_k] - (x - \tilde{\omega}) \quad E[\omega|\Omega_i] \leq x \leq E[\omega|\Omega_{i+1}], \quad (1.26)$$

where $i = 1, \dots, W$, and the 0-th segment is $x = 0$, $-\infty \leq x \leq E[\omega|\Omega_1]$.

$\hat{\mathcal{L}}_{lb}(x, \omega)$ and $\mathcal{L}_{lb}(x, \omega)$ are direct applications of Jensen's inequality. This section next presents upper bounds of $\hat{\mathcal{L}}(x, \omega)$ and $\mathcal{L}(x, \omega)$. A piecewise linear upper bound, i.e. Edmundson-Madanski's bound, can be obtained by shifting the lower bound up by a value e_W in Lemma 1.2.10 and Lemma 1.2.11.

Lemma 1.2.12. Consider the upper bound of $\hat{\mathcal{L}}(x, \omega)$.

$$\hat{\mathcal{L}}_{ub}(x, \omega) = \sum_{i=1}^W p_i \max(x - E[\omega|\Omega_i], 0) + e_W \quad (1.27)$$

is a piecewise linear function with $W + 1$ segments. The i -th segment of $\hat{\mathcal{L}}_{ub}(x, \omega)$ is

$$\hat{\mathcal{L}}_{ub}^i(x, \omega) = \sum_{k=1}^i p_k - \sum_{k=1}^i p_k E[\omega|\Omega_k] + e_W \quad E[\omega|\Omega_i] \leq x \leq E[\omega|\Omega_{i+1}], \quad (1.28)$$

where $i = 1, \dots, W$, and the 0-th segment is $x = e_W$, $-\infty \leq x \leq E[\omega|\Omega_1]$.

Lemma 1.2.13. The i -th segment of $\mathcal{L}_{ub}(x, \omega)$ can be written as, by applying Lemma 1.2.9,

$$\mathcal{L}_{ub}^i(x, \omega) = \sum_{k=1}^i p_k - \sum_{k=1}^i p_k E[\omega|\Omega_k] - (x - \tilde{\omega}) + e_W \quad E[\omega|\Omega_i] \leq x \leq E[\omega|\Omega_{i+1}], \quad (1.29)$$

where $i = 1, \dots, W$, and the 0-th segment is $x = e_W$, $-\infty \leq x \leq E[\omega|\Omega_1]$.

Having established the above results, one must decide how to partition the support Ω of ω in order to obtain good bounds. Rossi et al. (2015) proposed a simple and effective approach, which splits the support Ω into W disjoint regions with uniform probability mass, $p_i = Pr\{\omega \in \Omega_i\} = \frac{1}{W}$, $i = 1, \dots, W$. Once the probabilities p_i are fixed, $E[\omega|\Omega_i]$ is uniquely determined. The maximum approximation error e_W is attained at one of the breakpoints of its piecewise linear approximation.

Piecewise linear approximation in normal distribution. Consider a special case of the standard normally distributed random variable Z . The piecewise linear approximation of Z can be easily extended to the general case of a normally distributed variable ζ with mean μ and standard deviation σ via the following lemma.

Lemma 1.2.14. $\hat{\mathcal{L}}(x, \zeta)$ can be expressed in terms of the standard normal cumulative distribution function as

$$\hat{\mathcal{L}}(x, \zeta) = \sigma \int_{-\infty}^{\frac{x-\mu}{\sigma}} \Phi(x) dt = \sigma \hat{\mathcal{L}}\left(\frac{x-\mu}{\sigma}, Z\right). \quad (1.30)$$

Example. Consider a standard normal random variable Z , the support $\Omega = [-\infty, +\infty]$ of Z is partitioned into two segments. The probability mass $p_1 = 0.5$, the breakpoint $E[Z|\Omega_1] = 0$, and the maximum approximation error $e_W = \frac{1}{\sqrt{2\pi}}$ are obtained. This work presents Jensen's lower bounds and Edmundson-Madanski upper bounds for the loss function $\mathcal{L}(x, Z)$ in Fig. 1.13 and its complementary function $\hat{\mathcal{L}}(x, Z)$ in Fig. 1.14.

Note that the piecewise linear approximation parameters, i.e.: probability mass, breakpoints, and approximation errors, of $\hat{\mathcal{L}}_{lb}(x, Z)$ with up to eleven segments are reported in Rossi et al. (2014).

1.2.5 Mixed integer linear programming

This section introduces the Mixed Integer Linear Programming (MILP), in which the piecewise linear functions discussed in the last section can be readily applied. This dissertation extensively applies MILP models that are built upon these piecewise linear approximations to tackle a wide range of stochastic inventory problems. This section is mainly based on Wolsey (1998).




Figure 1.13: Two-segment piecewise linear lower bound $\mathcal{L}_{lb}(x, Z)$ and upper bound $\mathcal{L}_{ub}(x, Z)$ for the loss function $\mathcal{L}(x, Z)$




Figure 1.14: Two-segment piecewise linear lower bound $\hat{\mathcal{L}}_{lb}(x, Z)$ and upper bound $\hat{\mathcal{L}}_{ub}(x, Z)$ for the complementary of first order loss function $\hat{\mathcal{L}}(x, Z)$

The general formulation of an MILP is as follows.

$$\min c^T x \quad (1.31)$$

subject to,

$$Ax \leq b \quad (1.32)$$

$$x \geq 0 \quad (1.33)$$

$$x_i \in \mathbb{Z}, \forall i \in \mathcal{I}. \quad (1.34)$$

Note that c and b are vectors and A is a matrix. The objective function and all constraints are linear, and decision variables are integers or binaries.

MILP is widely used mainly because of the widespread availability of effective state-of-the-art MILP solvers, such as GUROBI, IBM-CPLEX, and XPRESS Optimizer, which incorporate many advanced techniques, e.g. (Jünger et al., 2009).

Branch-and-Bound. Most state-of-the-art MILP solvers are based on the branch-and-bound algorithm, which is an implicit enumeration method that uses a search tree to find an optimal solution. The general scheme is presented in Algorithm 1. There are three main steps: pick a variable and divide the problem into two subproblems at this variable (line 5-7); solve the LP-relaxation to determine the best possible objective function value for the node (line 8); prune the branch of the tree if either the subproblem is infeasible or the best achievable objective value is worse than a known optimum (line 10-15).

```

1 activeset := ∅;
2 bestval := NULL;
3 currentbest := NULL;
4 while activeset ≠ ∅ do
5   | choose a branching node  $k \in \textit{activeset}$ ;
6   | remove node  $k$  from activeset;
7   | generate the children of  $k$ , child  $i = 1, \dots, n_k$ ;
8   | generate optimistic bounds  $obj_i$  for each child  $i = 1, \dots, n_k$ ;
9   | for  $i = 1$  to  $n_k$  do
10  | | if  $obj_i$  is worse than bestval then
11  | | | kill child  $i$ ;
12  | | else if child  $i$  is a complete solution then
13  | | | bestval :=  $obj_i$ ;
14  | | | currentbest := child  $i$ ;
15  | | else
16  | | | add child  $i$  to activeset;
17  | | end
18  | end
19 end

```

Algorithm 1: The Branch-and-Bound algorithm

Big M. It is a very well-known modelling approach in MILP to use a binary variable δ to control whether linear constraint (1.32) is active or not. Then, this constraint can be reformulated as big- M , i.e., $Ax - b \leq M\delta$, M is a large enough value. However, the choice of values of M is nontrivial; a large M will result in

computational efficiency issue, while a small M may affect the solution quality. MILP models presented in this work come with advantage of being able to use the indicator constraints (Belotti et al., 2016), i.e., $\delta = 0 \Rightarrow Ax - b = 0$, which can avoid difficulties arising from the use of big- M .

Piecewise syntax. Consider the `piecewise` syntax in IBM-CPLEX, by means of which a piecewise function can be specified by giving a set of slopes which represent the linear variation for each linear segment, a set of breakpoints at which slopes change, and the function value at a known point (Leenaerts and Van Bokhoven, 2013).

```

piecewise(i in 1..W){
  slope[i] -> breakpoint[i];
  slope[W+1]
}(<knownpoint>,<valuepoint>)<value>;

```

Figure 1.15: The syntax of the `piecewise` command

The `piecewise` syntax is presented in Figure 1.15. W is the number of breakpoints of the piecewise function. `slope[i]` and `breakpoint[i]` denote slope and breakpoint of segment i . Segment i goes from breakpoint $(i - 1)$ to breakpoint (i) . `<valuepoint>` is the function value at a known point `<knownpoint>`. Finally, `<value>` represents the value at which we evaluate the function.

The `piecewise` syntax is adopted in Chapter 2 (Paper I) for computing parameters of the optimal inventory control policy.

1.3 Related works

This section discusses related works in stochastic inventory control and in particular on the (s, S) policy in Section 1.3.1, (R, S) policy in Section 1.3.2, and the Stochastic Joint Replenishment Problem in Section 1.3.3, which are extensively studied under different settings in this work.

Relevant background on deterministic and stochastic inventory control has been presented in Section 1.2.1. For a more in depth discussion, readers may refer to textbooks (Silver et al., 1998; Zipkin, 2000; Snyder and Shen, 2011). Since the pioneer study of Harris (1913) who proposed the EOQ model to answer the two fundamental questions of how much and how often to place an order, a vast body of literature has emerged on inventory control. For a thorough overview of inventory control literature, readers could refer to (Aggarwal, 1974; Yano and Lee, 1995; Ullah and Parveen, 2010; Glock et al., 2014; Bushuev et al., 2015). Most existing literature still presents applications to constant and dynamic deterministic demand; however, the study regarding stochastic demand has received increasing attention due to its practical relevance (Bastos et al., 2017). This work belongs to the growing literature on stochastic inventory control.

1.3.1 (s, S) policy

This section first discusses single-item stochastic inventory problems under (s, S) policy, and then briefly introduces literature on multi-item and multi-echelon stochastic inventory problems. A summary of literature surveyed in this section is presented in Table 1.3.1.

1.3.1.1 Single-item inventory system

This section surveys the literature on single-item (s, S) inventory systems under independent demand, service level, dependent demand, and robust demand settings.

Independent demand. The first study on proving the optimality of (s, S) policy could date back to Scarf (1960), which considered the finite horizon dynamic inventory system with fixed and linear ordering costs, and convex holding and penalty costs, by leveraging a novel property: K -convexity. Iglehart (1963) later extended the work of Scarf (1960) to the stationary infinite horizon problem. Veinott Jr (1966) further replaced Scarf's (Scarf, 1960) hypothesis that the one period expected costs are convex by a weaker assumption that the negatives of

Literature	demand	review	lead time	ordering cost	penalty cost	service level	lost sale	echelon	item	lifetime
Scarf (1960)	independent, stochastic	periodic	zero, fixed	fixed, unit	unit			single	single	
Iglehart (1963)	independent, stochastic	periodic	zero, fixed	fixed, unit	unit			single	single	
Venott Jr (1966)	independent, stochastic	periodic	fixed	fixed, unit	unit			single	single	
Ancja and Noori (1987)	independent, stochastic	periodic	zero	fixed, unit	unit			single	single	
Veinott Jr and Wagner (1965)	independent, stochastic	periodic	fixed	fixed, unit	unit			single	single	
Ehrhardt (1979)	independent, stochastic	periodic	fixed	fixed, unit	unit			single	single	
Askin (1981)	independent, stochastic	periodic	fixed	fixed, unit	unit			single	single	
Federguen and Zipkin (1984)	independent, stochastic	periodic	zero	fixed, unit	unit			single	single	
Zheng and Federguen (1991)	independent, stochastic	periodic	zero	fixed, unit	unit			single	single	
Bollapragada and Morton (1999)	independent, stochastic	periodic	zero	fixed, unit	unit			single	single	
Bashyam and Fu (1998)	independent, stochastic	periodic	stochastic	fixed, unit	unit	fill-rate		single	single	
Schneider and Ringuest (1990)	independent, stochastic	periodic	fixed	fixed	unit	γ		single	single	
Song and Zipkin (1993)	dependent, Markov-modulated Poisson	continuous	stochastic	fixed, unit	unit			single	single	
Sethi and Cheng (1997)	dependent, Markovian demand	periodic	zero	fixed, unit	unit	β		single	single	
Beyer and Sethi (1997)	dependent, Markovian demand	periodic	zero	fixed, unit	unit			single	single	
Chen and Song (2001)	dependent, Markov-modulated demand	periodic	fixed	unit	unit			multiple	single	
Lian et al. (2009)	dependent, Markovian renewal demand process	continuous	zero	fixed	unit			single	single	exponential
Hu et al. (2016)	dependent, Markov-modulated demand	periodic	zero	fixed, unit	unit			multiple	single	
Özekici and Parlari (1999)	dependent, Markov-modulated demand	periodic	zero	fixed, unit	unit			single	single	
Cheng and Sethi (1999)	dependent, Markov-modulated demand	periodic	zero	fixed, unit	unit		included	single	single	
Xu et al. (2010)	dependent, Erlang demand	periodic	zero	fixed, unit	unit		included	single	single	
Weiss (1980)	independent, Poisson	continuous	zero	fixed, unit	unit		included	single	single	fixed
Liu and Lian (1999)	independent, renewal demand process	continuous	zero	fixed	fixed, unit			single	single	fixed
Kalpakam and Sapna (1994)	independent, Poisson	continuous	zero	fixed, unit	unit		included	single	single	exponential
Liu and Yang (1999)	independent, Poisson	continuous	exponential	fixed, unit	unit			single	single	exponential
Ravichandran (1995)	independent, Poisson	continuous	fixed	fixed, unit	unit			single	single	exponential
Liu (1990)	independent, Poisson	continuous	zero	fixed	fixed, unit			single	single	exponential
Liu and Shi (1999)	independent, renewal demand process	continuous	zero	fixed	fixed, unit			single	single	exponential
Kalin (1980)	independent, stochastic	periodic	zero	fixed, unit	unit			single	single	exponential
Qiu et al. (2017)	robust	periodic	zero	fixed, unit	unit			multiple	single	
Cohen et al. (1988)	independent, stationary	periodic	fixed	fixed, unit	unit	fill-rate	included	multiple	multiple	

Table 1.1: A summary of literature surveyed on stochastic (s, S) policy

the one period expected costs are unimodal. Moreover, Aneja and Noori (1987) extended Scarf's results to a special case where the shortage cost associated with the failure to meet demand is a combination of a lump-sum cost, incurred whenever a stock-out is recognised and a proportional cost which is determined by the size of the shortage. While these early works proved the optimal policy is (s, S) for stochastic inventory systems, they did not present any computational study.

Veinott Jr and Wagner (1965) proposed a two-step method for computing optimal (s, S) policy parameters under assumptions of independent stochastic demand, constant lead time, fixed and linear ordering costs, and linear holding and penalty costs. The first step is to find the collection of all (s, S) policies that minimise the long-run average cost per unit time for some suitably small and fixed value of initial inventory level; then, to search the optimal (s, S) policies that minimises the long-run average cost per unit time for every initial inventory level.

Ehrhardt (1979) pointed out that the iterative methods for computing optimal policies in Veinott Jr and Wagner (1965) is computationally prohibitive for practical implementation, and requires the complete specification of the demand distribution. This is particularly unrealistic in practical settings. Therefore, the author presented the Power Approximation approach for computing the optimal (s, S) policy parameters, which is easy-to-implement, and requires as input only the mean and variance of demand over lead time. Computational experiments showed that the proposed Power Approximation approach yields expected total costs that typically are well within one percent of optimal.

Additionally, Federgruen and Zipkin (1984) proposed a computationally efficient policy iteration method for computing near-optimal (s, S) policies. Zheng and Federgruen (1991) proposed a search-based heuristic for approximating optimal (s, S) policy parameters built upon a number of new properties of the long-run average cost, a new upper bound for optimal order-up-to level and a new lower bound for the optimal reorder point. This algorithm is easy to understand and computationally efficient for both finite and infinite planning horizons. However, it can only solve stationary stochastic inventory control problems.

Askin (1981) extended the Silver-Meal algorithm and introduced a two-step heuristic for approximating nonstationary stochastic (s, S) policies. It first determines the order-up-to level S , and then calculates the reorder point s for which the minimised expected cost per unit period of placing an order is equal to the expected cost without placing an order. This heuristic is computationally faster, easier to implement and less sensitive to distant data. However, Bollapragada

and Morton (1999) pointed out that it is prohibitively expensive for other distributions except for normal, because it relies on the computation of convolutions of demand.

In contrast, Bollapragada and Morton (1999) proposed an efficient two-step stationary approximation approach for computing near-optimal (s, S) policies. The first step is to tabulate the values of reorder point, order-up-to level, the expected time between consecutive orders, and the total mean demand over the optimal expected time between two orders by solving the stationary problem for values of the mean demand at regular intervals to span all possible values of mean demand. Here, a large number of efficient algorithms exist for generating the stationary table, e.g. (Federgruen and Zipkin, 1984; Zheng and Federgruen, 1991). The second step is to read the optimal (s, S) policy parameters from the precomputed table by averaging nonstationary parameters (like mean demand) between two successive ordering periods. Computational experiments demonstrated that this approach is on average 400 times faster than dynamic programming, because it only involves calculating the review interval length followed by a precomputed table lookup. It does not depend on the system parameters. However, this stationary approximation approach assumes that only the expected demand varies across time periods. In a more general situation, more than one demand parameter might vary, which might lead to the computational complexity of this method, especially for the stationary table generation.

A recent comparison study (Dural-Selcuk et al., 2016) estimated the optimality gaps of Askin (1981) and Bollapragada and Morton (1999) at 3.9% and 4.9%, respectively; these figures are in line with those reported in the original works. Chapter 2 (Paper I) considerably improves these optimality gaps by means of a novel MILP-based heuristic for tackling this class of problems.

Even though nonstationary stochastic demand is more common in industrial settings with seasonal patterns, trend, businesses cycles, and limited-life items, stationary inventory control policies are more preferred in real-life settings. In this context, Tunc et al. (2011) investigated the cost of using a stationary policy as an approximation to the optimal nonstationary one. It adopted the cost-optimal (s, S) policy as a frame of reference, and compared the cost performance of the optimal nonstationary (s, S) policy with the best possible stationary (s, S) policy. It showed that stationary policies might be efficient approximations to optimal nonstationary policies when demand information contains high uncertainty, setup costs are high, and penalty costs are low.

Service level constraints. In spite of a vast body of literature that has been

developed on the cost minimisation problem, it has been widely recognised that penalty costs, and in particular, the cost of losing customer goodwill, are usually difficult to assess (Brown, 1967). On the contrary, service level measures are particularly popular in practice. Under this environment, the objective is to minimise the cost function, defined only in terms of fixed and unit ordering cost and holding costs, subjects to the constraint that the solution satisfies a prescribed level of customer service. In the literature, the “fill-rate” service level measure, the fraction of demand that is met directly from stock on hand, and the γ service level, the average backlog per period should not exceed a pre-defined level $(1 - \gamma)$ of the average demand per period, are commonly adopted.

The “fill-rate” service level measure was investigated by Bashyam and Fu (1998), which proposed a feasible directions procedure that is simulation-based for periodic review (s, S) inventory system with stochastic lead time and constraint of fill-rate service level measure. Computational experiments illustrated that the algorithm achieved 5% of optimality in 95% of the cases, and within 2% of optimality in 68% of the cases. Schneider and Ringuest (1990) proposed an analytic approximation, similar to Ehrhardt (1979), for (s, S) inventory systems with γ service level constraints. Computational results showed that the approximation gives a γ service level which is within 1% of the required service level in most cases.

Dependent demand. Classical inventory models have assumed that the inter-demand times are identically independent distributed. However, in real life, environmental factors, such as economic conditions and market conditions, can have a significant effect on demand (Song and Zipkin, 1993). It is therefore necessary to develop inventory models considering demand correlations.

Under the assumption of dependent demand, Markov processes have been widely used in stochastic inventory models with setup costs for modelling demand. Song and Zipkin (1993) presented a continuous-time, discrete-state formulation with a Markov-modulated Poisson demand and with linear costs of inventory and backlogging. It showed that the optimal policy is of state-dependent (s, S) type when the ordering cost consists of both a fixed cost and a linear cost. Sethi and Cheng (1997) generalised the work of Song and Zipkin (1993) to general demand, state-dependent convex inventory/backlog costs, and proved the optimality of (s, S) policy. Beyer and Sethi (1997) incorporated convex surplus cost into the model and proved the optimality from the viewpoint of minimising the long-run average cost of inventory/backlog and ordering. Similarly, Chen and Song (2001) proved the optimality of (s, S) policy from the viewpoint of minimising the long-

run average holding and backorder costs in the system.

Lian et al. (2009) proposed the first perishable inventory model with Markovian renewal demand, and proved the optimal policy is (s, S) type. Multi-echelon models incorporating Markov-modulated demand are discussed in (Chen and Song, 2001; Hu et al., 2016). Other studies have shown that the optimality of the (s, S) policy can be generalised to cases involving unbounded Markovian demand (Beyer et al., 1998), unreliable suppliers (Özekici and Parlar, 1999), and polynomial growth demand, returns, and cost functions (Li, 2013).

Additionally, unlike most studies in the literature only conducted under the assumption that unsatisfied demand is fully backlogged, (Cheng and Sethi, 1999; Xu et al., 2010) considered the lost sales situation is occurring in many retail establishments such as department stores and supermarkets. Cheng and Sethi (1999) use the analysis of the Markovian demand model with backlogging to analyse the lost sales case; in particular, they established the optimal (s, S) policy. Xu et al. (2010) further proved the (s, S) policy is optimal with Erlang distributed demand and lost sale settings.

Perishable items. Many commodities in practice may undergo deterioration in quality or functionality while they are in-storage or on-shelf and may have to be discarded eventually without being used. For example, in supermarkets, the fresh food may deteriorate gradually before it should get consumed; the electronic products may age while still in storage; fashion may become out of date when the seasons change. These perishable goods seize a large proportion of inventory so that the ordering policies determined by the conventional inventory model are not appropriate. This setting requires building up a particular perishable inventory model to study the structure of optimal ordering policies (Lian et al., 2009).

Weiss (1980) proved the optimality of the (s, S) -type policy for the continuous review perishable inventory system with zero lead time considering both backlogging and lost-sales. Liu and Lian (1999) relaxed the assumption of Poisson demand in Weiss (1980) to general renewal demand process. It showed that the expected total cost is monotone, convex, or unimodal in both the reorder point and the order-up-to level, and derived a closed-form for computing the optimal (s, S) policy parameters. Their computational study demonstrated that these properties are not affected by the coefficient of variation of the demand process. Hence, it is reasonable to believe that the (s, S) policy is also optimal for models with general renewal demand process.

Finite lifetime adds another dimension to inventory problems and often makes the modelling and analysis far more difficult than the corresponding inventory

models for items with infinite lifetime. As a result, replenishment lead time is usually not considered in existing literature such as Kalpakam and Arivarignan (1988). However, replenishment lead time is important in practice since it might significantly affect the replenishment decisions (Liu and Yang, 1999). Kalpakam and Sapna (1994) extended the work of Kalpakam and Arivarignan (1988) to a (s, S) system with exponential lead time and Poisson demand in which items deteriorate at a constant rate. They derived an exact long-run average cost expression under assumptions that unsatisfied demand is lost sales and the number of outstanding replenishment orders are restricted to at most one at any given time.

Similarly, Liu and Yang (1999) assumed an exponential lifetime and an exponential lead time (s, S) system. However, compared to existing literature, it allows backorders and does not limit the number of outstanding replenishment orders at any given time. Using the matrix analytic approach, they obtained analytical results on the performance measures and developed a numerical optimisation procedure.

Moreover, Ravichandran (1995) showed that the perishable inventory model of (s, S) type with realistic assumptions, like non-instantaneous lead time, constant lifetime, and Poisson demand process, is complex and intractable. They presented an expression representing the long-run average cost per unit time in operating under the (s, S) policy, which may be used as an objective to identify the optimal reorder level for a specified ageing phenomenon. The closed-form expression allows a numerical search procedure to obtain optimal parameters.

Liu (1990) considered a continuous review, exponential life time (s, S) inventory system in which demand occurs in single units following a Poisson process. Liu and Shi (1999) provided a comprehensive treatment of the model in Liu (1990) with a general renewal demand process. They presented a simple but important relation of two system performance measures, the expected inventory level and cycle time, with whom the total cost function can be easily constructed. By means of the supplementary variable method, the authors built a Markov process so as to obtain the expected length of the cycle time.

Lian et al. (2009) further extended the work of Liu and Shi (1999) to the continuous review, exponential life time (s, S) inventory system with a general Markovian renewal demand process. They presented the first perishable inventory model with Markovian renewal demand and derived the analytical expression for the expected recycle time, the expected total cost rate function, and the optimal ordering policies.

Robust demand. In the literature, many researchers have made great effort to identify effective inventory policies to determine when and how much to order of a product. Establishing an effective inventory policy often requires an in-depth analysis of the nature of the target business. Traditional inventory models, particularly for a multi-period setting, usually assume that the demand distribution of a product and all of its parameters are completely known. However, these assumptions may not hold in many practical situations. Thus, the solutions based on such assumptions may lead to severe constraint violations even under very small perturbations (Beyer and Sendhoff, 2007). As pointed out in Bertsimas and Thiele (2006), an optimal inventory policy heavily tuned to a particular demand distribution may perform very poorly for another demand distribution bearing the same uncertainty parameter.

Therefore, Qiu et al. (2017) investigated a finite-horizon single-product periodic-review inventory system with uncertainty in demand probability distributions. They proved that the (s, S) policy is optimal for nonstationary distribution-free inventory problems, and proposed a dynamic robust model for approximating optimal policy parameters for the box and the ellipsoid uncertainty sets, which can be transformed into tractable linear and second-order cone programmes. However, the proposed models and solution approaches are only validated for the boxes, and the ellipsoid uncertainty sets.

1.3.1.2 Multi-item inventory system

The (s, S) policy is also adopted in the multi-item periodic-review stochastic inventory system. It is represented by (σ, \vec{S}) since multiple items are involved, where $\sigma \subset \mathcal{R}^N$, and $\vec{S} \in \mathcal{R}^N$. Note that the definition of σ is general; its shape and properties are literately unknown. There is no guarantee of σ by convex, or even connected. Johnson (1967) characterised the optimal policy for the stationary case and introduced the (σ, \vec{S}) policy, in which $\sigma \subset \mathcal{R}^N$ and $\vec{S} \in \mathcal{R}^N$; in this policy one orders up to \vec{S} if the inventory level \vec{I} is in set σ and $\vec{I} \leq \vec{S}$, otherwise one does not order. Kalin (1980) showed that, when $\vec{I} \in \sigma$ and $\vec{I} \not\leq \vec{S}$, there exists $\vec{S}(\vec{I}) \geq \vec{I}$ such that the optimal policy is to order up to $\vec{S}(\vec{I})$, this policy is named as $(\sigma, \vec{S}(\cdot))$ policy. Ohno et al. (1994) proposed an algorithm for computing an optimal ordering policy $(\sigma, \vec{S}(\cdot))$ for a periodic view multi-item inventory system. Ohno and Ishigaki (2001) further proposed a policy iteration method to compute an exact optimal policy by leaving properties of the optimal policy for continuous-time inventory problems with compound Poisson demand. Gallego and Sethi (2005) gave the general definition of K -convexity in

\mathcal{R}^N , which encompasses both the joint ordering and individual ordering case; it derived an optimal policy for the two-item deterministic inventory problem with a joint ordering cost. However, the computation of the optimal (σ, \vec{S}) policies is still a difficult task.

Due to the complexity of the (σ, \vec{S}) policy, literature on the computation of the optimal policy parameters is still lacking. Chapter 4 (Paper III) presents models can be used to identify whether a given initial inventory level \vec{I}_0 is in set σ .

1.3.1.3 Multi-echelon inventory system

Cohen et al. (1988) considered the multi-echelon inventory system which is characterised by the use of emergency shortage shipments, since demand not met at each stocking location is passed up to higher echelons, and can be viewed as a lost sales at a local stocking location. Therefore, demand at a particular location is composed of normal replenishment orders coming from lower echelons, and emergency shipment orders from both lower echelons and local stocking locations. They introduced an approximate model for the (s, S) inventory system with the lost sale setting, which is the first model that distinguishes between classes of prioritised demand under the assumptions, and developed an efficient and effective solution heuristic for solving the fill-rate service level measure optimisation problem. The computational study indicated good performance which deteriorates as the fill rate requirement and lead time increase.

1.3.2 (R, S) policy

This section surveys literature on the stochastic (R, S) policy. It first presents works on computing stochastic (R, S) policies with service level constraints, backorders, and the combination of service level constraints and backorders, and then on the multiple sourcing (R, S) inventory system. A summary of the literature surveyed in this section is presented in Table 1.2.

Although the (s, S) policy is proved to be optimal for both stationary and nonstationary stochastic demand under certain conditions (Scarf, 1960; Iglehart, 1963; Veinott Jr, 1966; Aneja and Noori, 1987), it performs poorly with respect to the system nervousness resulting from revisions in the original plan which in turn result in different replenishment decisions in successive planning cycles (De Kok et al., 1997; Heisig, 1998, 2001). In this regard, the (R, S) policy, in which the timing and the order-up-to level are fixed at the beginning of the planning

Literature	review	lead time	ordering cost	penalty cost	service level	lost sale	lifetime
Bookbinder and Tan (1988)	periodic	zero	fixed		α		infinite
Tarim and Kingsman (2004)	periodic	zero	fixed, unit		α		infinite
Rossi et al. (2008)	periodic	zero	fixed		α		infinite
Tarim and Smith (2008)	periodic	zero	fixed		α		infinite
Tempelmeier (2007)	periodic	zero	fixed		α_p, β		infinite
Tarim et al. (2011)	periodic	zero	fixed		α		infinite
Tunc et al. (2014)	periodic	zero	fixed		α		infinite
Rossi et al. (2010)	periodic	stochastic	fixed		α		infinite
Rossi et al. (2012a)	periodic	stochastic	fixed		α		infinite
Pauls-Worm et al. (2014)	periodic	stochastic	fixed, unit		α		fixed
Tarim and Kingsman (2006)	periodic	zero	fixed, unit	unit			infinite
Rossi et al. (2012b)	periodic	zero	fixed, unit	unit			infinite
Hua et al. (2009)	rolling horizon	fixed	fixed	unit			infinite
Özen et al. (2012)	periodic	zero	fixed	unit	α		infinite
Rossi et al. (2015)	periodic	zero	fixed, unit	unit	α, β, β^c	included	infinite
Tunc et al. (2018)	periodic	zero	fixed	unit	α, β, β^c	included	infinite

Table 1.2: A summary of literature surveyed on single-item stochastic (R, S) policy

horizon, and the actual order quantities are determined after realised demand, provides a means of dampening the system nervousness (Silver et al., 1998). Kilic and Tarim (2011) revealed that the (R, S) policy has the potential to replace the (s, S) policy, especially for systems characterised by a low degree of flexibility to setup changes. Furthermore, Tunc et al. (2013) pointed out that the (R, S) policy performs nearly as good as the (s, S) policy, and it is an effective policy for coordinating supply chain inventories especially when setup-oriented system nervousness is of concern.

To find the optimal (R, S) policy parameters is known to be a difficult problem. As a result, the literature on computing policy parameters was not well-established until the recent three decades.

Service level constraints. Bookbinder and Tan (1988) is probably the first work for computing nonstationary stochastic optimal (R, S) policy parameters. It proposed a two-stage method for computing the minimised expected total cost comprising fixed ordering cost and holding cost under the α service level constraint. It firstly computes the timing of replenishments at the beginning of the planning horizon by using the Wagner-Whitin algorithm (Wagner and Whitin (1958)); then, determines buffer stock levels by solving a linear programming model.

Bookbinder and Tan (1988) ignored the unit ordering cost since it is assumed to be a constant, and has no effect on the determination of the best schedule. However, Tarim and Kingsman (2004) showed that the unit ordering cost could not be neglected since it will affect the objective function. It formulated a mixed integer programming (MIP) model for simultaneously determining the timing of replenishment and the corresponding order-up-to levels. However, this paper does not address computational performance issues.

Since the work of Tarim and Kingsman (2004), considerable attention has been given to address the computational performance under the same assumption. Rossi et al. (2008) suggested a novel concept, global chance-constraints, based on which they proposed an exact stochastic constraint programming approach for computing (R, S) policy parameters. However, the number of binary variables increases polynomial in the number of periods. In contrast, Tarim and Smith (2008) proposed the constraint programming formulation which reduces the number of decision variables. The computational study showed that the constraint programming formulation is more solvable than the MIP model. Two independent domain reduction methods were proposed to improve the computational performance of the MIP model and the constraint programming formulation.

Instead of considering the α service level as discussed in (Bookbinder and Tan, 1988; Tarim and Kingsman, 2004; Rossi et al., 2008; Tarim and Smith, 2008), Tempelmeier (2007) claimed that this criterion does not provide much information about the performance of an inventory system as it may be zero although the system works almost perfectly. Thus, Tempelmeier (2007) proposed a new formulation for computing the optimal (R, S) policies under the α_p service level constraint, the expected proportion of order cycles with no stock-outs, and the β service level constraint, the expected proportion of demand routinely filled from stock.

Tarim et al. (2011) relaxed the original MIP model in Tarim and Kingsman (2004) to the shortest path problem. Numerical experiments demonstrated that feasible solutions are obtained in the majority of test instances; in the case of infeasibility, the solution can be used to generate a feasible one, which provides an upper bound. A simple branch-and-bound procedure that implements the relaxation approach at each node of the search tree is used to search for an optimal solution. Numerical evidence shows that these bounds are tight, leading to an efficient and fast search procedure. Tunc et al. (2014) reformulated the original MIP model as a deterministic equivalent MIP model, analytically verified the linear relaxation of the reformulation is stronger, and numerically showed the computational efficiency of the reformulation.

Above-mentioned works assume zero supplier lead time. However, the lead time uncertainty, which in various industries is an inherent part of the business environment, has a detrimental effect on inventory systems. For this reason, there exist two inventory control studies analysing the impact of supplier lead time uncertainty on the (R, S) policy under the α service level constraints. Rossi et al. (2010) proposed a stochastic constraint programming model to address this issue. This is the first work in the literature. Rossi et al. (2012a) developed two constraint-based local search methods based on a coordinate descent strategy for finding the near-optimal (R, S) policy parameters under nonstationary stochastic demand and lead time.

Additionally, Pauls-Worm et al. (2014) considered the single-item single-stock location inventory control problem for a perishable item under the α service-level constraint. They first formulated the stochastic programming model, and then proposed an MILP approximation originally introduced in Tarim and Kingsman (2004) for computing the optimal (R, S) policies for nonperishable items. Computational experiments demonstrated that the MILP model guarantees that in 96.4% of the periods the service level requirements are met with an error tolerance

of 1%.

Penalty costs. Tarim and Kingsman (2006) addressed the single-item, non-stationary stochastic (R, S) inventory control problem under the assumption of backorders being allowed. They presented an MILP model for computing the optimal policy parameters by adopting a piecewise linear approximation to the nonlinear terms in the cost function. They derived an explicit formulation for a special case of the standard normally distributed demand.

The power of the MILP model of Tarim and Kingsman (2006) is the ability to be reused for any normally distributed demand; however, to evaluate the accuracy of the piecewise linear approximation technique is not easy. Therefore, Rossi et al. (2012b) provided an exact constraint programming formulation for computing the optimal (R, S) policies, and proposed a dedicated cost-based filtering method to improve the performance of the search. The proposed solution approach can be further used to gauge the solution quality of the MILP model in Tarim and Kingsman (2006).

Hua et al. (2009) further proposed a static-dynamic uncertainty model under the rolling horizon setting. It determines the timing of replenishment and the corresponding order-up-to levels over the planning horizon but implements only the decisions of the first period. It then uses the rolling horizon approach in the next period when the inventory level and the demand distribution are updated based on the demand realisation.

Service level constraints and penalty costs. There are few attempts for building models to incorporate both service level constraints and penalty costs. Özen et al. (2012) developed a mathematical model and dynamic programming based solution algorithm for computing near-optimal (R, S) policy parameters. To solve large-scale problems, an approximation heuristic is proposed to approximate the true cost function, and a relaxation heuristic is proposed to relax a constraint in the original problem. Computational experiments illustrated that both heuristics perform well in terms of solution quality and computation time.

Rossi et al. (2015) generalised the above-discussed models on computing (R, S) policy parameters, and presented a unified MILP model which applies both when the unmet demand is backordered and in lost sale settings and accommodates variants of service level constraints such as α , β and β_{cyc} (the expected fraction of demand that is routinely satisfied from stock for each replenishment cycle). Computational studies demonstrated the effectiveness and flexibility of the proposed approach.

Recently, Tunc et al. (2018) presented an extended MIP model that blends heuristic methods originally introduced by Tunc et al. (2014) and Rossi et al. (2015). As a result, this formulation features the computational efficiency of Tunc et al. (2014) and the modelling variety of Rossi et al. (2015). The proposed formulation is essentially designed to approximate the original nonlinear cost function with a prior piecewise linear function. Nonetheless, it also developed a dynamic cut generation approach to deploy the model with a non-prior approximation of the function while guaranteeing an arbitrary level of precision.

Multiple sourcing. The (R, S) policy is also adopted in the multiple sourcing inventory system. Janssen and de Kok (1999) considered an inventory system with two suppliers. General supply agreements are made with the main supplier to deliver a fixed quantity Q at every review period. An order is placed at the second supplier such that the inventory position is raised to the order-up-to position S if the inventory position at the beginning of the review period is below the order-up-to position S . An algorithm for the determination of the decision parameters Q and S was developed such that the long-run expected average costs per time unit (the summation of the holding, and fixed and unit ordering costs) are minimised subject to a service level constraint. The computational results illustrated the effectiveness and profitability of the multiple sourcing strategy above the single sourcing strategy.

1.3.3 Joint replenishment problem

The problem of controlling inventory of a multi-item system under joint replenishment has been receiving considerable attention for the past several decades. For a thorough review of literature readers could refer to (Silver and Peterson, 1985; Goyal and Satir, 1989; Van Eijs et al., 1992; Khouja and Goyal, 2008; Bastos et al., 2017). This section focuses on the literature for tackling stochastic JRPs. In particular, control policies that have been extensively studied, e.g.: (s, c, S) policy (Balintfy, 1964), (R, T) policy (Atkins and Iyogun, 1988), (Q, S) policy (Renberg and Planche, 1967), (s, S) policy (Kalpakam and Arivarignan, 1993), $P(s, S)$ policy (Viswanathan, 1997), $Q(s, S)$ policy (Nielsen and Larsen, 2005), and (Q, S, T) policy (Özkaya et al., 2006). A detailed discussion on these policies can be found in Chapter 4 (paper III). However, this dissertation adopts the existing (R, S) policy, originally introduced in Bookbinder and Tan (1988) for tackling single-time stochastic inventory problems, to tackle stochastic multi-item inventory problems subject to joint ordering costs. Additionally, this work differs

from existing literature in being able to solve stationary as well as nonstationary demand.

1.4 Thesis statement

This section summarises works described in this dissertation in Section 1.4.1, and then highlights contributions of this dissertation in Section 1.4.2, followed by the respective summary for each of the following chapters.

1.4.1 Summary

This work mainly concentrates on the application of mathematical programming based models for a wide range of problems in the field of stochastic inventory control.

An important problem in inventory control is the determination of near-optimal inventory control policies under nonstationary stochastic demand for the classical setting captured in Scarf (1960). As discussed in Section 1.2 and 1.3, this problem is of importance in business practice, and significant research has emerged in recent decades. We consider different existing formulations of this problem under modelling assumptions of independent demand, correlated demand, and joint replenishment settings. For these formulations, existing approaches in the literature have revealed three drawbacks.

Firstly, existing approaches for computing near optimal policy parameters for the classical setting captured in Scarf (1960) provide fairly wide optimality gaps. A recent study Dural-Selcuk et al. (2016) shows that the only two approaches available in the literature (Askin, 1981; Bollapragada and Morton, 1999) provide fairly loose optimality gaps approximately 3.9% and 4.9%.

Secondly, most existing approaches are not easily implementable. Due to the combinatorial nature of stochastic inventory control, numerical approaches proposed in the literature are search-based (Viswanathan, 1997; Bollapragada and Morton, 1999; Özkaya et al., 2006), which results in considerable effort in computer coding. Although some easy-to-implement methods have been presented in (Rossi et al., 2015; Tunc et al., 2018), most literature still requires tedious computer coding.

Thirdly, no approach in the literature relaxes the assumption of independence under nonstationary demand. Models hereby presented close this gap, and are also able to compute near-optimal replenishment policies under a wide range of time-series-based demand processes.

Finally, even though various policies exist in the literature for solving stochastic joint replenishment problems, these works only consider stationary demand.

In order to overcome drawbacks of existing approaches in the literature, novel,

near-optimal, easy-implementabl, and broadly applicable mathematical programming heuristics are discussed in the following chapters. These approaches are built upon stochastic programming bounding techniques, i.e., Jensen’s lower bound and Edmundson-Madanski upper bound, and can relax restrictive assumptions in the stochastic inventory control literature thus addressing unexplored settings.

These models are used to compute near-optimal policy parameters under the classical setting in Scarf (1960), they are also used to relax the classical assumption of independent demand and investigate the case in which demand is correlated, including a collection of time-series-based demand processes. Finally, they are used to compute near optimal control policies for a nonstationary multi-item inventory system.

Unlike other existing approaches in the literature, these methods can be implemented and solved by using off-the-shelf mathematical programming packages, such as IBM ILOG optimisation studio and XPRESS Optimizer, and do not require tedious computer coding.

Computational studies on the approximation of nonstationary (s, S) policy parameters showed that the proposed models provide the best optimality gaps (generally below 0.3%) in the literature for the single-item single-location stochastic inventory problem. By applying the same model for tackling correlated demand, computational experiments demonstrated a tight optimality gap (2.28%). When considering stochastic joint replenishment problems, these models produced competitive optimality gaps against existing policies in the literature.

1.4.2 Contributions

This work contributes to the literature of stochastic inventory control by applying these novel, near-optimal, and easy-to-implement models to a wide range of problems.

- **Application 1.** Chapter 2 applies new mathematical programming based models to compute the near optimal inventory control policy parameters for single-item single-stock location stochastic inventory problems under the classical assumption of independent demand. This is the first MILP-based heuristic for computing near-optimal (s, S) policies since it was initially proved to be optimal for this class of problems by Scarf (1960). This chapter also combines the previously introduced MILP model and a binary search procedure to solve large-scale problems. Computational experiments demonstrate that both models yield the best optimality gap, i.e., generally

below 0.3%, and guarantee reasonable computational time.

- **Application 2.** Chapter 3 relaxes the long-standing assumption of independent demand in the literature of stochastic inventory control and investigates the case in which demand is correlated, including a collection of time-series-based demand processes. This chapter develops a stochastic model which captures the (R, S) policy under correlated demand. This is the first time the (R, S) policy has been expressed in the form of a functional equation. This chapter also presents an MILP-based heuristic for approximating optimal (R, S) policies under normally distributed demand featuring correlation across periods as well as under a collection of time-series-based demand processes. Computational experiments show that the optimality gap is 2.28%, and the average computational time is acceptable.
- **Application 3.** Chapter 4 introduces an MILP model to approximate the inventory control policy parameters of joint replenishment under nonstationary demand. This is the first work to tackle nonstationary joint replenishment problems in the literature. This modelling method can also be used to approximate the optimal control rule for this class of problems, known as (σ, \vec{S}) policy. A computational study shows that the proposed method is competitive in terms of cost performance. Most importantly, this model has the advantage of being able to tackle nonstationary joint replenishment problems which have not been addressed in the literature.

1.4.3 Chapter 2 (Paper I): Computing non-stationary (s, S) policies using mixed integer linear programming

This paper is a joint work with Roberto Rossi, Belen Martin-Barragan, and S. Armagan Tarim.

The computation of nonstationary (s, S) policies is a long-standing problem in the literature of stochastic inventory control. Since the pioneering work of Scarf (1960) which proved the optimality of (s, S) policies for a class of dynamic inventory models, several approaches were proposed to compute the optimal policy parameters.

In the literature, studies on this topic can be categorised into stationary and nonstationary. A number of studies investigated the computation of stationary (s, S) policy parameters, such as (Federgruen and Zipkin, 1984; Zheng and Federgruen, 1991; Feng and Xiao, 2000), whereas only two approaches (Askin, 1981;

Bollapragada and Morton, 1999) exist for approximating nonstationary (s, S) policies. A recent comparative research Dural-Selcuk et al. (2016) estimated the optimality gap of these two approaches at 3.9% and 4.9%, respectively.

In order to find simple and yet effective heuristic methods for computing (s, S) policy parameters, this work presents two MILP-based heuristics, which leverage two key building blocks: a modelling technique initially discussed in Rossi et al. (2015) (presented in Section 1.2.4) and K -convexity of the problem cost function originally discussed in Scarf (1960) (presented in Section 1.2.1.2).

First, an MINLP-based heuristic is built, which can be linearized via the approach originally discussed in Rossi et al. (2014) and can be implemented in OPL by adopting a `piecewise` expression (presented in Section 1.2.4). The resulting mathematical programming models can be solved by using off-the-shelf optimisation packages, such as IBM ILOG OPL studio and XPRESS Optimizer.

To tackle larger instances, we then combine the previously introduced MINLP model and a binary search procedure; this latter approach requires dedicated codes but scales better than the previous one.

Extensive computational experiments are conducted to demonstrate the computational performance of the proposed MILP-based heuristics. This paper first investigates the performance of both models by contrasting costs of the policy obtained with our models against costs of the optimal policy obtained via stochastic dynamic programming. We observe optimality gaps are generally below 0.3%. We then assess the computational performance of the binary search based heuristics. Computational experiments show that the computational efficiency is reasonable: around 748.20 seconds on average.

In this work, my contribution can be summarised as follow. I survey the literature related to the long-standing problem of the computation of single-item single-location nonstationary (s, S) policies under independent demand. I adapt the results in Rossi et al. (2015) (piecewise linear approximations) and in Scarf (1960) (K -convexity) to implement the MILP model and the binary search strategy in IBM ILOG optimisation studio. I also implement the benchmark stochastic dynamic programming in Matlab. I conduct numerical studies on small instances and large-scale instances both including 540 instances. I organise all material and write the paper.

1.4.4 Chapter 3 (Paper II): (R, S) policy with correlated demand

This paper is a joint work with Roberto Rossi, Belen Martin-Barragan, and S.

A long-standing assumption in existing studies is that random demand in each period is independent of demand in other periods. However, as Song and Zipkin (1993) pointed out, environmental factors, such as economic conditions, market conditions, and any exogenous condition, have significant effects on the demand for a product, the supply, and the cost structure. This paper relaxes the assumption of independent demand in Chapter 2 (Paper I).

Correlated demand has been previously investigated in the inventory literature. Authors attempted to either prove the optimality of the (s, S) policy or compute optimal policy parameters with different types of demand correlations over the planning horizon. However, to the best of our knowledge, no study on computing (R, S) policies under correlated demand and time-series-based demand processes exists.

This paper considers a periodic-review single-item single-stock location lot-sizing problem under non-stationary stochastic correlated demand. We develop a stochastic model which captures the (R, S) policy under correlated demand. Note that it is also the first time that the (R, S) policy has been expressed in the form of a functional equation.

We then leverage properties of conditional distributions and present an MILP-based heuristic for approximating optimal (R, S) policies under normally distributed demand featuring correlation across periods as well as under a collection of time-series-based demand processes. Our approach offers a stable replenishment plan while effectively hedging against uncertainty. Our model can be easily implemented and solved by using off-the-shelf mathematical programming packages such as IBM ILOG optimisation studio.

This work further illustrates how to adapt the model to a collection of time-series-based demand processes: the autoregressive (AR) process, the moving average (MA) process, the autoregressive moving-average (ARMA) process, and the autoregressive with autoregressive conditional heteroskedasticity (AR-ARCH) process.

An extensive computational study is carried out to investigate the performance of the proposed MILP heuristic. We first assess the behaviour of the optimality gap and the computational efficiency of the MILP heuristic on multivariate normally distributed demand. Computational experiments show that the optimality gap is 2.28%, and the average computational time is 0.1s. We then assess the computational performance of the MILP model on time-series-based demand processes. Computational experiments demonstrate that the average

computational time is 0.68s.

The computational study demonstrate that the proposed MILP model is computationally efficient and accurate. Moreover, in contrast to existing approaches in the literature, it can tackle higher order time-series-based demand processes.

In this work, my contribution can be summarised as follow. I survey the literature related to the computation of near-optimal policies under correlated demand. I implement the stochastic dynamic programming under correlated demand in Matlab, then implement the MILP-based model for approximating optimal (R, S) policies under normally distributed demand featuring correlation across periods, and then extend this model to cover a wide range of time-series-based demand processes including AR, MA, ARMA, and AR-ARCH processes. Moreover, I implement these models under receding horizon control settings in Java with the help of Prof. Rossi. I conduct computational experiments on the test bed including 432 instances. I organise all material and write this paper; during the writing process Dr Martin-Barragan provides valuable help in the derivation of the (R, S) policy in the form of a functional equation.

1.4.5 Chapter 4 (Paper III): Nonstationary (R, S) policies for joint replenishment inventory systems

This paper is a joint work with Roberto Rossi, Belen Martin-Barragan, and S. Armagan Tarim.

This paper extends Chapter 3 (Paper II) to a multi-item nonstationary stochastic inventory system under the joint replenishment setting, which has been receiving increasing attention over the past decades. A number of policies has been presented for tackling stochastic joint replenishment problems, e.g.: the (Q, S, T) policy (Özkaya et al., 2006), $Q(s, S)$ policy (Nielsen and Larsen, 2005), $P(s, S)$ policy (Viswanathan, 1997), (Q, S) policy (Pantumsinchai, 1992), MP policy (Atkins and Iyogun, 1988), $(s, c, S)_M$ policy (Melchior, 2002), and $(s, c, S)_F$ policy (Federgruen et al., 1984).

This work considers the stochastic joint replenishment problem under the (R, S) policy. We present an MILP model, built upon the piecewise linear approximation technique originally introduced in Rossi et al. (2014) (presented in Section 1.2.4) for approximating the optimal policy parameters. In contrast to existing methods, this model can be easily implemented and solved by using existing off-the-shelf optimisation packages.

We then demonstrate that our MILP model can be used to approximate the

optimal control rule for stochastic joint replenishment problems, known as (σ, \vec{S}) policy. Due to the complexity of σ , it is impractical to derive a closed form expression for it. However, our MILP model can be used to determine whether given initial inventory levels $\vec{I}_0 \in \sigma$. This allows a decision maker to decide if it is optimal or not to trigger a joint replenishment order at a given period.

We assess the cost performance of the (R, S) policy by comparing its cost performance against existing policies on data sets of Atkins and Iyogun (1988) and Viswanathan (1997). We notice that both data sets contain some unusual lot sizing instances; more specifically, instances for which the group as well as item fixed ordering costs become negligible in comparison to holding costs. To focus on meaningful lot sizing instances, we filter test instances of both data sets by using the following conditions: $K > b \geq h$, which are broadly accepted in the lot sizing literature. We also check the order frequency in each period, and we discard instances in which joint orders for all items are issued too frequently, and for which the optimal policy essentially degenerates to a base stock policy.

Computational experiments on the data set of (Atkins and Iyogun, 1988) demonstrate that the (R, S) policy fully dominates other competing policies in the literature in 2 out of 10 test instances considered. The (R, S) policy performs better than the $Q(s, S)$, (Q, S) , MP , $(s, c, S)_M$, and $(s, c, S)_F$ policies with average cost improvement of 0.07%, 1.74%, 0.89%, 2.84%, and 7.02%, respectively; however, the (Q, S, T) and $P(s, S)$ policies perform slightly better than the (R, S) policy with average cost improvement of 0.09% and 0.17%.

Computational experiments on the data set of (Viswanathan, 1997) illustrate that the (R, S) is the best policy in 13 out of 31 test instances. The (R, S) performs better than the $Q(s, S)$, $P(s, S)$, (Q, S) , MP , and $(s, c, S)_F$ policies with average cost improvement of 0.37%, 0.37%, 1.81%, 1.71%, and 1.67%; while (Q, S, T) policy performs slightly better than them with an average cost improvement of 0.19%.

Extensive computational studies indicate that although the (R, S) policy does not fully dominate all existing policies, it comes with the additional advantage of being able to tackle stationary and nonstationary demand. To the best of our knowledge, this is the first work on computing (R, S) policies for nonstationary joint replenishment problems.

In this work, my contribution can be summarised as follow. I survey the literature on tackling the joint replenishment problem. I implement the MILP model on computing near-optimal (R, S) policy parameters in IBM ILOG optimisation studio and the shortest path formulation in Matlab and IBM ILOG optimisation

studio. I conduct computational experiments on both data sets of Atkins and Iyogun (1988) and Viswanathan (1997). I organise all material and write the paper.

1.5 Future work

This section first summarises which questions remain open for each of the following chapters, and then describes potential research areas where the mathematical programming based models can be successfully applied.

Chapter 2. This chapter presents the first MILP-based model for computing optimal (s, S) policy parameters, which provides the best optimality gaps in the literature, and can be implemented and solved without tedious computer coding. However, this MILP-base model only can effectively tackle small-size instances. To preserve the advantage of relying on an MILP model, it is useful to investigate efficient reformulations, valid inequalities, or to explore cut generation techniques that enhance computational performance. We believe that the investigation of efficient reformulations and techniques to improve the computational performance of the proposed MILP-based models is a promising research direction.

Chapter 3. This chapter proposes the first MILP model for computing the optimal (R, S) policies under normally distributed demand featuring correlation across periods as well as under a collection of time-series-based demand processes. It is possible to explore the cost of using an independent demand policy as an approximation to the optimal dependent demand policy. A future research direction is to adopt the (R, S) policy as a frame of reference, and compare the cost performance of the (R, S) policy under independent demand with the best possible (R, S) policy under correlated demand.

Chapter 4. This chapter proposes an MILP model for tackling stochastic joint replenishment problem under an (R, S) policy. Although this model does not dominate all existing policies in the literature, it comes with the additional advantage of being able to tackle nonstationary demand which has not been addressed in the literature. It is worthwhile to investigate the cost performance of the proposed method under a rolling horizon setting.

All presented mathematical programming based models are novel, near-optimal, and easy-to-implement; most importantly, they share a common underpinning modelling strategy. The application of this modelling strategy to other problems in the field of stochastic inventory control is a promising area to investigate.

A first direction worth investigating is stochastic lead time. As discussed in Section 1.3.1, only Bashyam and Fu (1998) proposed a simulation-based method and Song and Zipkin (1993) proposed a search-based algorithm for computing optimal (s, S) policies. As discussed in Section 1.3.2, (Rossi et al., 2010, 2012a) introduced constraint programming models for approximating the optimal (R, S)

policies. It would be interesting to investigate if the near-optimal and easy-to-implement mathematical programming models presented in this work can be adapted to tackle stochastic lead time.

Secondly, many commodities may undergo deterioration in quality or functionality while they are in-storage or on-shelf and may have to be discarded eventually without being used (Lian et al., 2009). In this setting, optimal ordering policies determined by the conventional (nonperishable) inventory models are not appropriate. As discussed in Section 1.3.1, a number of studies have been done to either prove the optimality of the (s, S) policy under perishable items or compute near-optimal policy parameters. Additionally, as discussed in Section 1.3.2, Pauls-Worm et al. (2014) proposed an MILP model for computing the near-optimal (R, S) policies. However, this model might only work well for a very high in-stock probability. Also in this case it would be interesting to investigate if the near-optimal and easy-to-implement mathematical programming models presented in this work can be adapted; a first step in this direction has been discussed in Gutierrez-Alcoba et al. (2016).

Thirdly, it is natural to integrate stochastic inventory control with other business operations, such as carbon emissions. As pointed out in Bushuev et al. (2015), the literature on this aspect has been scattered, and the EOQ and Newsvendor are the most frequently used models for sustainable lot sizing. The integration of other business operations is therefore a promising area to which one may apply the modelling strategy that underpins all the mathematical programming models discussed in this work.

1.6 Conclusions

This work presents novel, near-optimal, and easy-to-implement mathematical programming based models that are built upon the application of stochastic programming bounding techniques: Jensen's lower bound and Edmundson-Madansky upper bound. It applies these modelling methods to a wide range of problems in the field of stochastic inventory control under nonstationary demand.

A long-standing problem in stochastic inventory control is the approximation of nonstationary (s, S) policy parameters for the single-item inventory problem under the assumption of independent demand. Chapter 2 (Paper I) presents the first MILP-based heuristic for computing near-optimal policies. To tackle larger instances, it combines the MILP model and a binary search procedure. Extensive computational studies illustrated that these MILP-based heuristics yielded an average optimality gap 0.03% which is the best in the literature.

Chapter 3 (Paper II) relaxes the independence assumption and investigates the case in which demand is correlated. This is the first time that a stochastic model has been developed, which captures the (R, S) policy under correlated demand. An MILP model is presented for approximating optimal (R, S) policies under normally distributed demand featuring correlation across periods as well as under a collection of time-series-based demand processes. Computational experiment shows that the optimality gap is 2.28% and the average computational time is acceptable.

Finally, Chapter 4 (Paper III) tackles nonstationary joint replenishment problems under (R, S) policy. Computational experiments show that the proposed model is competitive in terms of cost performance with other state-of-the-art policies. It fully dominates existing policies in 2 out of 9 test instances on the data set of (Atkins and Iyogun, 1988) and in 13 out of 31 test instances on the data set of (Viswanathan, 1997). Although it does not fully dominate all existing competitor strategies, it comes with the advantage of being able to tackle nonstationary demand which has not been addressed in the literature.

Extensive computational studies show that the proposed models are easily implemented and solved by using existing off-the-shelf mathematical packages. Most importantly, they do not require tedious computer coding.

In summary, stochastic inventory control is an active research area in Operational Research. This work contributed to the literature of stochastic inventory control by tackling a wide range of problems under unexplored settings using near-optimal and easy-to-implement mathematical programming heuristics that

are built upon stochastic bounding techniques. Future contributions to the literature on stochastic inventory control may focus on improving these modelling methods, applying these models to further relax long-standing assumptions in the literature such as fixed lead time and fixed lifetime, and integrating with other research areas such as routing problems and carbon emissions.

Chapter 2

Paper I: Computing non-stationary (s, S) policies using mixed integer linear programming

Abstract

This paper addresses the single-item single-stock location non-stationary stochastic lot sizing problem under the (s, S) control policy. We first present a mixed integer non-linear programming (MINLP) formulation for determining near-optimal (s, S) policy parameters. To tackle larger instances, we then combine the previously introduced MINLP model and a binary search approach. These models can be reformulated as mixed integer linear programming (MILP) models which can be easily implemented and solved by using off-the-shelf optimisation software. Computational experiments demonstrate that optimality gaps of these models are less than 0.3% of the optimal policy cost and computational times are reasonable.

2.1 Introduction

Stochastic lot sizing is an important research area in inventory theory. One of the landmark studies is Scarf (1960), which proved the optimality of (s, S) policies for a class of dynamic inventory models. The (s, S) policy features two control

parameters: s and S . Under this policy, the decision maker checks the opening inventory level at the beginning of each time period: if it drops to or below the reorder point s , then a replenishment should be placed to reach the order-up-to level S . Unfortunately, computing optimal (s, S) policy parameters remains a computationally intensive task.

Since Scarf's landmark study, the (s, S) policy has been object of extensive research. For instance, (Johnson and Thompson, 1975; Sethi and Cheng, 1997; Chen and Song, 2001; Hu et al., 2016) investigated demand correlation; more recently, (Qiu et al., 2017; Lim and Wang, 2017) investigated demand distributional ambiguity.

In the literature, studies on (s, S) policy can be categorized into stationary and non-stationary. A number of studies investigated the computation of stationary (s, S) policy parameters, e.g. (Iglehart, 1963; Veinott Jr and Wagner, 1965; Archibald and Silver, 1978; Stidham, 1977; Sahin, 1982; Federgruen and Zipkin, 1984; Zheng and Federgruen, 1991; Feng and Xiao, 2000). However, there has been an increasing recognition that lot-sizing studies need to be undertaken for non-stationary environments (Graves, 1999).

In this work, we focus on the single-item single-stock location stochastic lot-sizing problem under non-stationary demand, fixed and unit ordering cost, holding cost and penalty cost. Only two studies investigated computations of (s, S) policy under non-stationary stochastic demand (Askin, 1981; Bollapragada and Morton, 1999).

Askin (1981) adopted the "least cost per unit time" approach in selecting order-up-to levels and reorder points under a penalty cost scheme. Decision makers first determine desired cycle lengths and order-up-to levels. Then, reorder points are decided by means of a trade-off analysis between expected costs per period in cases of ordering and not ordering.

As Bollapragada and Morton (1999) pointed out, the approach discussed by Askin (1981) is computationally expensive because of the need of convolving demand distributions. In contrast, Bollapragada and Morton (1999) proposed a stationary approximation heuristic for computing optimal (s, S) policy parameters. Firstly, decision makers precompute pairs of (s, S) values for various demand parameters and tabulate results. Here, a large number of efficient algorithms exist for generating the stationary table, e.g. (Federgruen and Zipkin, 1984; Zheng and Federgruen, 1991; Feng and Xiao, 2000). Secondly, order-up-to levels and reorder points can be read from stationary tables by averaging the demand parameters over an estimate of the expected time between two orders. However,

this algorithm relies upon complex code, particularly for generating stationary tables.

(Askin, 1981; Bollapragada and Morton, 1999) do not provide a satisfactory solution to the problem of computing near-optimal (s, S) policy parameters: they rely on ad-hoc computer coding and provide relatively large optimality gaps. A recent computational study (Dural-Selcuk et al., 2016) estimated the optimality gap of (Askin, 1981; Bollapragada and Morton, 1999) at 3.9% and 4.9%, respectively; these figures are in line with those reported in the original works. These drawbacks motivate the investigation of simple and yet effective heuristic methods for computing (s, S) policy parameters; methods that do not need dedicated computer coding and that can provide better optimality gaps.

The aim of this paper is to introduce two new heuristics to compute near-optimal (s, S) policy parameters. We build upon Rossi et al. (2015), which discussed mixed-integer linear programming (MILP) heuristics for approximating optimal (R, S) policy parameters — under this policy, the replenishment intervals R and order-up-to levels S are determined at the beginning of the planning horizon, while associated order quantities are decided only when orders are issued. The (R, S) policy is effective in dealing with system nervousness (Tunc et al., 2013), while the (s, S) policy is cost-optimal (Scarf, 1960). Our two mixed-integer nonlinear programming (MINLP)-based heuristics leverage two key building blocks: modeling techniques originally discussed in Rossi et al. (2015), and K -convexity of the problem cost function, originally discussed in Scarf (1960). In contrast to other approaches in the literature, our heuristics can be easily implemented and solved by using off-the-shelf mathematical programming packages such as IBM ILOG optimisation studio.

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

- We introduce the first mixed integer non-linear programming (MINLP) model to compute near-optimal (s, S) policy parameters.
- We show that this model can be approximated as a mixed integer linear programming (MILP) model by piecewise linearising the cost function; this approximation can be solved by using off-the-shelf software.
- To tackle larger instances, we combine the previously introduced MINLP model and a binary search procedure; this latter approach requires dedicated code, but scales better than the previous one.
- Computational experiments demonstrate that optimality gaps of our models are tighter than existing algorithms (Askin, 1981; Bollapragada and Morton,

1999) in the literature, and computational times of our models are reasonable.

The rest of this paper is organised as follows. Section 2.2 describes the problem setting and a stochastic dynamic programming (SDP) formulation. Section 2.3 discusses the notion of K -convexity and introduces relevant K -convex cost functions which are approximated by an MINLP model in Section 2.4. Section 2.5 presents an MINLP heuristic for approximating (s, S) policy parameters. Section 2.6 introduces an alternative binary search approach for computing (s, S) policy parameters. A detailed computational study is given in Section 2.7. Finally, we draw conclusions in Section 2.8.

2.2 Problem description

We consider a single-item single-stock location inventory management system over a T -period planning horizon. We assume that orders are placed at the beginning of each time period, and delivered instantaneously. Ordering costs $c(\cdot)$ comprise a fixed ordering cost K for placing an order, and a linear ordering cost c proportional to order quantity Q . Demand d_t in each period $t = 1, \dots, T$ is a independent random variable with known probability distribution. At the end of period t , a linear holding cost h is charged on every unit carried from one period to the next; and a linear penalty cost b is occurred for each unmet demand at the end of each time period.

For a given period $t = 1, \dots, T$, let I_{t-1} denote the opening inventory level and Q_t represent the order quantity.

The immediate expected holding and penalty costs at period t can be expressed as

$$f_t(I_{t-1}, Q_t) = E[h \max(I_{t-1} + Q_t - d_t, 0) + b \max(d_t - I_{t-1} - Q_t, 0)], \quad (2.1)$$

where “E” denotes the expectation taken with respect to the random demand d_t .

The ordering cost $c(Q_t)$ is defined as:

$$c(Q_t) = \begin{cases} K + c Q_t, & Q_t > 0 \\ 0, & Q_t = 0 \end{cases}$$

Let $C_t(I_{t-1})$ represent the expected total cost of an optimal policy over periods t, \dots, T when the initial inventory level at the beginning of period t is I_{t-1} . We

model the problem as a stochastic dynamic program (Bellman, 1966) via the following functional equation

$$C_t(I_{t-1}) = \min_{Q_t} \{c(Q_t) + f_t(I_{t-1}, Q_t) + E[C_{t+1}(I_{t-1} + Q_t - d_t)]\} \quad (2.2)$$

where

$$C_T(I_{T-1}) = \min_{Q_T} \{c(Q_T) + f_T(I_{T-1}, Q_T)\}$$

represents the boundary condition.

2.3 The optimality of (s, S) policies in stochastic lot sizing

Scarf (1960) proved that the optimal policy in the dynamic inventory problem is always of the (s, S) type based on a study of the function $G_t(y) + cy$, where

$$G_t(y) = f_t(y) + E[C_{t+1}(y - d_t)], \quad (2.3)$$

and y is the stock level immediately after purchases are delivered (see Scarf, 1960, Eq. (4)).

Since we consider a non-stationary environment, values of the (s, S) policy parameters will depend on the given period t . Let (s_t, S_t) denote the policy parameters for period t . Function $G_t(y) + cy$ can be used to identify optimal policy parameters (s_t, S_t) . In particular, the order-up-to level S_t is defined as the value minimising $G_t(y) + cy$; whereas the parameter s_t is given by the value $s_t < S_t$ such that $G_t(s_t) + cs_t = G_t(S_t) + cS_t + K$ (see Scarf, 1960, Eq. (5)). K -convexity of the function $G_t(y) + cy$ ensures the uniqueness of s_t and S_t (Scarf, 1960).

Example. We illustrate the concepts introduced on a 4-period example. Demand d_t is normally distributed in each period t with mean $\mu_t = 20, 40, 60, 40$, for $t = 1, \dots, 4$ respectively. Standard deviation σ_t of demand in period t is equal to $0.25\mu_t$. Other parameters are $K = 100$, $h = 1$, $b = 10$, and $c = 0$. We plot $G_1(y)$ in Fig. 2.1 for initial inventory levels $y \in [0, 200]$. Note that this problem is solved via the stochastic dynamic programming model presented in Section 2.2. The order-up-to level is $S_1 = 70$, $G_1(S_1) = 263$, the reorder point is $s_1 = 14$, and $G_1(s_1) = 363$. Note that $G_1(s_1) + cs_1 = G_1(S_1) + cS_1 + K$. The optimal policy is to order to 70 if the initial inventory drops below 14.




Figure 2.1: Plot of $G_1(y)$

2.4 MINLP approximation of $G_t(y)$ function

In this section, we exploit an MINLP model to approximate the function $G_t(y)$ in Eq. (2.3). Our model follows the control policy known as “static-dynamic uncertainty” strategy, known as (R, S) policy, originally introduced in Bookbinder and Tan (1988). Under this strategy, the timing of orders and order-up-to levels are expected to be determined at the beginning of the planning horizon, while associated order quantities are decided upon only when orders are issued. As illustrated in Rossi et al. (2015), this strategy provides a cost performance which is close to the optimal “dynamic uncertainty” strategy. However, optimal (s, S) parameters cannot be immediately derived from existing mathematical programming models operating under a static-dynamic uncertainty strategy, such as (Graves, 1999; Rossi et al., 2015). We next illustrate how a model operating under a static-dynamic uncertainty strategy can be used to approximate the function $G_t(y)$ in Eq. (2.3). In the rest of this section, without loss of generality, we focus on the case $G_1(y)$.

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

Let $t = 1, \dots, T$ and consider three sets of decision variables: \tilde{I}_t , the expected closing inventory level at the end of period t , with I_0 denoting the initial inventory level; δ_t , a binary variable which is set to one if an order is placed in period t ;

P_{jt} , a binary variable 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 in Constraints (2.9) and (2.10). Let \tilde{d}_{jt} denote the expected value of the demand over periods j, \dots, t , i.e. $\tilde{d}_{jt} = \tilde{d}_j + \dots + \tilde{d}_t$. Decision variables $H_t \geq 0$ and $B_t \geq 0$ represent end of period t expected excess inventory and back-orders, respectively. An MINLP formulation for the non-stationary stochastic lot-sizing problem under the “static-dynamic” uncertainty strategy, obtained following the modelling strategy in Rossi et al. (2015), is shown in Figure 2.2.

$$\min \left(\sum_{t=1}^T (K\delta_t + hH_t + bB_t) + c\tilde{I}_T + c \sum_{t=1}^T \tilde{d}_t - cI_0 \right) \quad (2.4)$$

Subject to, $t = 1, 2, \dots, T$

$$\delta_t = 0 \rightarrow \tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1} = 0 \quad (2.5)$$

$$\tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1} \geq 0 \quad (2.6)$$

$$\sum_{j=1}^t P_{jt} = 1 \quad (2.7)$$

$$P_{jt} \geq \delta_j - \sum_{k=j+1}^t \delta_k, \quad j = 1, 2, \dots, t \quad (2.8)$$

$$P_{jt} = 1 \rightarrow H_t = \hat{L}(\tilde{I}_t + \tilde{d}_{jt}, d_{jt}), \quad j = 1, 2, \dots, t \quad (2.9)$$

$$P_{jt} = 1 \rightarrow B_t = L(\tilde{I}_t + \tilde{d}_{jt}, d_{jt}), \quad j = 1, 2, \dots, t \quad (2.10)$$

$$P_{jt} \in \{0, 1\}, \quad j = 1, 2, \dots, t \quad (2.11)$$

$$\delta_t \in \{0, 1\} \quad (2.12)$$

Figure 2.2: The formulation of the non-stationary stochastic lot-sizing problem

The objective function (2.4) minimizes the expected total cost over the planning horizon. In the objective function, expected variable ordering costs are reformulated via $c \sum_{t=1}^T Q_t = c\tilde{I}_T + c \sum_{t=1}^T \tilde{d}_t - cI_0$ by using the reformulation strategy originally introduced in Tarim and Kingsman (2004) at p. 112 — note that term $c \sum_{t=1}^T \tilde{d}_t - cI_0$ is a constant. Constraints (2.5) are indicator constraints (Belotti et al., 2016) capturing the reorder condition. Constraints (2.6) are the inventory balance equations. Constraints (2.7) indicate the most recent replenishment before period t was issued in period j . Constraints (2.8) identify uniquely the period in which the most recent replenishment prior to t took place. Constraints (2.9)

and (2.10) are indicator constraints modelling end of period t expected excess inventory and back-orders by means of the first order loss function.

We now discuss how to adapt the model in Fig. 2.2 in order to compute, for a given y , an approximate value of $G_1(y) + cy$; see Eq. (2.3). We call this modified model MINLP- s , and use superscript s to label decision variables in this model.

In addition to constraints in Fig. 2.2, MINLP- s features constraint

$$\delta_1^s = 0, \quad (2.13)$$

which forces the model not to issue an order in period 1. When $\delta_1^s = 0$, the objective function (2.4) becomes

$$G_1^s(I_0^s) = \underbrace{hH_1^s + bB_1^s}_{f_1(I_0^s)} + \underbrace{\sum_{t=2}^T (K\delta_t^s + hH_t^s + bB_t^s)}_{\substack{\text{fixed ordering, holding, and} \\ \text{penalty cost for } t = \{2, \dots, T\}}} + \underbrace{c\tilde{I}_T^s + c \sum_{t=2}^T \tilde{d}_t - c\tilde{I}_1^s}_{\substack{\text{proportional ordering cost} \\ \text{for } t = \{2, \dots, T\}}} \quad (2.14)$$

\approx
 $E[C_2(I_0^s - d_1)]$

which denotes the expected total cost of controlling the system optimally over the planning horizon $1, \dots, T$ when the initial inventory level is I_0^s and no order is issued in period 1; hence $c(\tilde{I}_1^s + \tilde{d}_1 - I_0^s) = 0$.

MINLP- s can be approximated as an MILP model by using the approach discussed in Rossi et al. (2015) to piecewise linearise loss functions in constraints (2.9) and (2.10). For further details please refer to Appendix 2.A.

Example. In Fig. 2.3, we plot $G_1^s(y)$ obtained via the MILP- s for the same 4-period numerical example in Fig. 2.1 with initial inventory level $I_0^s \in [0, 200]$.

Since $G_1^s(y)$ approximates $G_1(y)$, we can now use $G_1^s(y) + cy$ to find approximate values \hat{S}_1 and \hat{s}_1 for S_1 and s_1 .

2.5 An MINLP-based model to approximate (s, S) policy parameters

In this section we exploit the results presented in the previous section to introduce an MINLP-based heuristic for approximating optimal (s, S) policies. To the best of our knowledge, this is the first MINLP model in the literature for computing near-optimal (s, S) policy parameters.

In a similar fashion to “MINLP- s ”, we introduce “MINLP- S ” to be the ap-




Figure 2.3: Plot of $G_1^s(y)$

proximation of $C_t(I_{t-1})$ in Eq. (2.2). Similarly to Eq. (2.14), let the objective function $C_1^S(I_0)$ of MINLP- S denote the expected total cost of controlling the system optimally over the planning horizon $1, \dots, T$ given the initial inventory level I_0 . We use the superscript S to represent decision variables in MINLP- S ,

$$C_1^S(I_0) = \sum_{t=1}^T (K\delta_t^S + hH_t^S + bB_t^S) + c\tilde{I}_T^S + c \sum_{t=1}^T \tilde{d}_t - cI_0. \quad (2.15)$$

MINLP- S imposes the constraint

$$\delta_1^S = 1, \quad (2.16)$$

which forces the model to place a replenishment in period 1.

In the MINLP- S model, \hat{S}_1 denotes an approximation of the optimal order-up-to level S_1 . Since $G_1^s(I_0^s)$ is an approximation of $G_1(I_0^s)$, by leveraging Scarf's result (see Scarf, 1960, Eq. (4)) on the study of $G(y) + cy$, we can identify $\hat{s}_1 = I_0^s$ such that $G_1^s(I_0^s) + cI_0^s = G_1^s(\hat{S}_1) + c\hat{S}_1 + K$. Therefore, we can approximate s_1 by imposing the constraint

$$G_1^s(I_0^s) + cI_0^s = C_1^S(I_0^S) + cI_0^S, \quad (2.17)$$

in which I_0^S represents an approximation \hat{S}_1 of the optimal order-up-to level S_1 ¹. Note that $C_1^S(I_0^S)$ includes the fixed ordering cost K because of Constraint (2.16);

¹ I_0^S , which is a dummy variable, should not be confused with the actual initial inventory level I_0 , which is needed to account for variable ordering costs.

variable ordering cost in $C_1^S(I_0^S)$ is zero because I_0^S is its global minimizer. Therefore, Eq. (2.17) is equivalent to $G_1^s(I_0^s) + cI_0^s = G_1^s(\hat{S}_1) + c\hat{S}_1 + K$.

Finally, since $s_1 \leq S_1$, we introduce an additional constraint to ensure that the reorder point is not greater than the order-up-to level,

$$I_0^s \leq I_0^S. \quad (2.18)$$

Note that, in contrast to the true value $G_1(y) + cy$, there is no guarantee that K -convexity holds for its approximation $G_1^s(y) + cy$. For some instances we may therefore have multiple values \hat{s}_1 such that Eq. (2.17) holds. As we will demonstrate in our computational study, leaving to the solver the freedom to choose one of such values in a non-deterministic fashion leads to alternative results.

MINLP- S and MINLP- s are connected by Eq. (2.17), in such a way that the order-up-to level S_1 and the reorder point s_1 are approximated simultaneously. For the joint MINLP model, in addition to decision variables in MINLP- S and MINLP- s , we consider I_0^S , a dummy variable representing the approximate order-up-to level \hat{S}_1 ; and I_0^s , which captures the approximate reorder point \hat{s}_1 .

Our holistic MINLP model objective features two parts: the first part, $C_1^S(I_0)$, comes from MINLP- S ; the second part, $G_1^s(I_0^s) + cI_0^s - f_1(I_0^s) \approx \mathbb{E}[C_2(I_0^s - d_1)]$, from MINLP- s . Note that the term $f_1(I_0^s)$, which enhances computational performance of the model, can be introduced because holding and penalty costs in period 1 for model MINLP- s are already uniquely determined by Eq. (2.17). After dropping the constant term $c \sum_{t=1}^T \tilde{d}_t - cI_0$ in the first part and the constant term $c \sum_{t=1}^T \tilde{d}_t$ in the second part, we minimise the following holistic objective function

$$\min \left(\sum_{t=1}^T (K\delta_t^S + hH_t^S + bB_t^S) + c\tilde{I}_T^S + \sum_{t=2}^T (K\delta_t^s + hH_t^s + bB_t^s) + c\tilde{I}_T^s \right); \quad (2.19)$$

Constraints of the joint MINLP model are those of both MINLP- S and MINLP- s in addition to the linking constraints (2.13), (2.16)-(2.18). After solving the joint MINLP model over planning horizon k, \dots, T , the estimated order-up-to level \hat{S}_k is equal to I_{k-1}^S , and the estimated reorder point \hat{s}_k is equal to I_{k-1}^s . As previously discussed, the joint MINLP model can be linearised via the piecewise-linear approximation proposed in Rossi et al. (2014). In our MILP model, (2.9) and (2.10) are modelled via the piecewise OPL expression (IBM, 2011). For a complete overview of the MILP model refer to Appendix 2.B.

Example. We now use the same 4-period numerical example in Fig. 2.3 to

demonstrate the modelling strategy behind the joint MINLP heuristic (MP). We observe that, for period 1, the approximated order-up-to level is $S_1 = 70.3$, the reorder point is $s_1 = 15.0$, and $G_1^s(s_1) = 366$ (363, after simulation) as shown in Fig. 2.1. By solving the joint MINLP repeatedly, s_t , S_t , and $G_t^s(s_t)$, for $t = 1, \dots, 4$, are estimated as shown in Table 2.1. We also compare our results against the optimal solutions obtained via SDP in Table 2.1; note that although different order-up-to levels, e.g. S_2 , are obtained, the optimal expected total costs are similar.

	MP				SDP			
t	1	2	3	4	1	2	3	4
s_t	15.0	29.0	58.1	29.0	14.0	29.0	58.0	28.0
S_t	70.3	54.0	117	54.0	70.0	141	114	53.0
$G_t^s(s_t)$	366	311	193	118	363	303	190	118

Table 2.1: Optimal (s, S) policy parameters obtained via the joint MINLP heuristic and the stochastic dynamic programming

2.6 A binary search approach to approximate (s, S) policy parameters

The joint MINLP heuristic presented in the last section is valuable, since it can be easily linearised into an MILP model that can be solved by off-the-shelf solvers. However, according to our experience, it can only effectively tackle small-size instances. To preserve the advantage of relying on an MILP model, one may investigate efficient reformulations, valid inequalities, or may explore cut generation techniques that enhance computational performances; we however choose to leave this investigation as future work.

In order to tackle larger-size problems, in this section we introduce an efficient approach that combines the model MINLP- s discussed in Section 2.5 and a binary search strategy. This approach relies on the MINLP models previously introduced, but it has the disadvantage of requiring dedicated code for the search procedure.

Our binary search strategy (Algorithm 2) is structured as follows.

Computation of S (line 2-3). We first let I_0^s to be a decision variable in MINLP- s and minimise $G_1^s(I_0^s) + cI_0^s$ to estimate the order-up-to level \hat{S}_1 .

Computation of s (line 5-17) Since $G_1^s(I_0^s)$ is an approximation of $G_1(y)$, we can conduct a binary search to approximate the reorder point \hat{s}_1 by $I_0^s \leq \hat{S}_1$

at which $G_1^s(I_0^s) + cI_0^s = G_1^s(\hat{S}_1) + c\hat{S}_1 + K$. When the binary search terminates, the estimated reorder point \hat{s}_k is equal to I_{k-1}^s .

By repeating this procedure (line 1) over the planning horizon k, \dots, T , we find pairs of \hat{S}_k and \hat{s}_k , where $k = 1, \dots, T$.

<p>Data: costs (ordering cost, holding cost, penalty cost), mean demand and standard deviation of each period, stepsize</p> <p>Result: pairs of s and S for each period</p> <pre> 1 for $k = 1$ to T do 2 Minimize MINLP-s in Section 2.5 in OPL; 3 Obtain $G_k^s(\hat{S}_k)$ and \hat{S}_k; 4 $low =$ a large negative integer; $high = \hat{S}_k$; 5 while $low < high$ do 6 $mid = low + \text{round}((high - low)/2)$; 7 Run the MINLP-$s$ with $I_{k-1}^s = mid$ in OPL; 8 Obtain current cost $G_k^s(I_{k-1}^s)$; 9 if $G_k^s(I_{k-1}^s) - G_k^s(\hat{S}_k) - K - c(\hat{S}_k - I_{k-1}^s) < 0$ then 10 $high = mid - \text{stepsize}$; 11 else if $G_k^s(I_{k-1}^s) - G_k^s(\hat{S}_k) - K - c(\hat{S}_k - I_{k-1}^s) > 0$ then 12 $low = mid + \text{stepsize}$; 13 else 14 $\hat{s}_k = I_{k-1}^s$; 15 $low = high$; 16 end 17 end 18 end </pre>

Algorithm 2: The binary search algorithm

Example. We illustrate the solution method discussed via the same 4-period numerical example presented in Fig. 2.1. We assume the step size of the binary search is 0.01. The order-up-to level $\hat{S}_1 = 70.3$ and $G_1^s(70.3) = 266$. We then set $low = -200$, $high = 70.3$. While $low < high$, the mid is updated via the comparison of $G_1^s(I_0^s) + K$ and $G_1^s(\hat{S}_1) + K$. Eventually, we obtain the reorder point $\hat{s}_1 = 15$ at which $G_1^s(\hat{s}_1) + c\hat{s}_1 = G_1^s(\hat{S}_1) + c\hat{S}_1 + K$. By repeating this procedure we obtain \hat{S}_t , \hat{s}_t , and $G_t^s(s_t)$, for each period $t = 1, \dots, 4$ as displayed in Table 2.2. After simulation, we obtain the expected total cost 363.

t	1	2	3	4
s_t	15.0	29.0	58.1	29.0
S_t	70.3	54.0	116	54.0
$G_t^s(s_t)$	366	311	193	118

Table 2.2: Near-optimal (s, S) policy parameters obtained via the binary search approach

2.7 Computational experiments

In this section we present an extensive analysis of the heuristics discussed in Sections 2.5 (MP) and 2.6 (BS). We first design a test bed featuring instances defined over an 8-period planning horizon (Section 2.7.1). On this test bed, we assess the behaviour of the optimality gap and the computational efficiency of both the MP and BS heuristics. Then we assess the computational performance of the BS heuristics on a test bed featuring larger instances on a 25-period planning horizon (Section 2.7.2). For all cases, MINLP models are solved by employing the piecewise linearisation strategy discussed in Rossi et al. (2015), which can be easily implemented in OPL by means of the `piecewise` syntax. Numerical experiments are conducted by using IBM ILOG CPLEX Optimization Studio 12.7 and MATLAB R2014a on a 3.2GHz Intel(R) Core(TM) with 8GB of RAM.

2.7.1 An 8-period test bed

We consider a test bed which includes 540 instances. Specifically, we incorporate ten demand patterns displayed in Fig. 2.4. These patterns comprising two life cycle patterns (LCY1 and LCY2), two sinusoidal patterns (SIN1 and SIN2), a stationary pattern (STA), a random pattern (RAND), and four empirical patterns (EMP1, ..., EMP4). Full details on the experimental set-up are given in Appendix 2.C. Fixed ordering cost K ranges in $\{200, 300, 400\}$, proportional ordering cost c ranges in $\{0, 1\}$, and the penalty cost b takes values in $\{5, 10, 20\}$. We assume that demand d_t in each period t is independent and normally distributed with mean μ_t and coefficient of variation $c_v \in \{0.1, 0.2, 0.3\}$; note that $\sigma_t = c_v \mu_t$. Since we operate under the assumption of normality, our models can be readily linearised by using the piecewise linearisation parameters available in Rossi et al. (2014). However, the reader should note that our proposed modelling strategy is distribution independent, see Rossi et al. (2015).

We set the SDP model discussed in Section 2.2 as a benchmark. We compare against this benchmark in terms of optimality gap and computational time. First




Figure 2.4: Demand patterns in our computational analysis

of all, we obtain optimal parameters for each test instance by implementing an SDP algorithm in MATLAB. Then, we solve each instance by adopting both modelling heuristics presented in Section 2.5 and 2.6. Specifically, for the MP heuristic we employ seven segments in the piecewise-linear approximations of B_t and H_t (for $t = 1, \dots, T$) in order to guarantee reasonable computational performances; for the BS heuristic, whose computational performance is only marginally affected by an increased number of segments in the linearisation, we employ eleven segments and a step size 0.1. To estimate the cost of the policies obtained via our heuristics, we simulate all policies via Monte Carlo Simulation (10,000 replications).

Table 2.3 gives an overview of optimality gaps (%) of methods discussed in this study for different pivoting parameters. It is difficult to make a general remark with respect to demand pattern, and fixed ordering cost; while the proportional ordering cost has a negative correlation with the optimality gap. An increase in proportional ordering cost slightly reduces the optimality gap. While an increase in penalty cost increases the optimality gap. Specifically, when penalty cost increases from 10 to 20, the optimal gap rises from 0.25% to 0.42% and from 0.27% to 0.35%, respectively. Similarly, an increase in coefficient of variation increases the optimality gap. For example, the optimality gap of the BS heuristic increases significantly from 0.16% to 0.40% as the coefficient of variation increases from 0.1 to 0.3. Overall, the average optimality gap of the MP heuristic is 0.29%, and that of the BS heuristic is 0.26%. This discrepancy ought to be expected, since in the case of the BS method a higher number of segments has been employed.

Existing heuristics Askin (1981) and Bollapragada and Morton (1999) were reimplemented by Dural-Selcuk et al. (2016) and assessed on a test bed that neatly resembles the one adopted in this work. As shown in Dural-Selcuk et al. (2016), Askin's optimality gap is 3.9%, and Bollapragada and Morton's is 4.9%. The optimality gap of our heuristic is 0.29% when seven segments are employed in the piecewise linearisation, and it drops to 0.26% when eleven segments are employed. Our models outperform both Askin (1981) and Bollapragada and Morton (1999) in terms of optimality gap on the test bed here considered.

We also assess the accuracy of our models by comparing the costs predicted by our models against the costs obtained via simulation. We note that both MP and BS heuristics have high accuracy for the 8-period numerical experiments. For further details please refer to Table 2.8 in Appendix 2.D.

Table 2.4 shows computational times of our models for different pivoting parameters. Note "STDEV" in Table 2.4 represents the standard deviation. Overall,

Modelling methods	MP	BS
Demand pattern		
LCY1	0.25	0.33
LCY2	0.11	0.18
SIN1	0.13	0.20
SIN2	0.10	0.19
STA	0.50	0.14
RAND	0.16	0.22
EMP1	0.41	0.35
EMP2	0.86	0.52
EMP3	0.15	0.19
EMP4	0.28	0.28
Fixed ordering cost		
200	0.31	0.29
300	0.24	0.22
400	0.34	0.27
Proportional ordering cost		
0	0.33	0.29
1	0.26	0.23
Penalty cost		
5	0.21	0.16
10	0.25	0.27
20	0.42	0.35
Coefficient of variation		
0.1	0.22	0.16
0.2	0.26	0.22
0.3	0.40	0.40
Average	0.29	0.26

Table 2.3: Average optimality gaps % of the 8-period numerical experiment for different pivoting parameters

the computational time of BS method remains stable for different set-up parameters; while that of MP and SDP algorithms fluctuate. We observe that the fixed ordering cost, proportional ordering cost, penalty cost, and coefficient of variation do not have significant effect on the computational efficiency of BS and SDP algorithms. However, the computational time of MP heuristic drops significantly with the increase of fixed ordering cost, and proportional ordering cost; while grows greatly with the increase of the coefficient of variation. On average, the computational time of MP, BS, and SDP are 7.89, 7.01, and 53.03 seconds.

Settings	MP		BS		SDP	
	Mean	STDEV	Mean	STDEV	Mean	STDEV
Demand pattern						
LCY1	3.54	0.98	7.17	1.21	13.58	0.86
LCY2	6.26	4.52	7.29	1.06	13.61	0.81
SIN1	4.67	3.20	6.48	0.69	13.31	1.11
SIN2	4.15	1.89	6.41	0.630	13.60	0.82
STA	5.52	3.68	6.48	0.72	9.95	2.29
RAND	3.60	0.87	7.11	1.32	710.12	2.95
EMP1	7.65	6.21	7.32	0.96	121.81	28.60
EMP2	14.03	13.60	7.28	1.19	107.20	7.37
EMP3	14.32	11.81	7.02	0.83	104.41	10.17
EMP4	15.12	15.35	7.52	1.20	122.71	27.79
Fixed ordering cost						
200	10.29	11.18	7.11	1.00	53.03	51.94
300	7.17	6.94	6.99	1.00	53.07	51.98
400	6.19	5.40	6.93	1.08	52.99	51.91
Proportional ordering cost						
0	8.49	9.06	7.64	0.99	60.21	60.12
1	7.28	7.57	6.38	0.59	45.85	40.85
Penalty cost						
5	8.05	7.92	6.96	0.90	53.08	52.03
10	8.72	10.60	6.86	1.02	52.74	51.94
20	6.84	8.84	7.17	1.14	52.97	51.85
Coefficient of variation						
0.1	6.42	6.16	7.00	1.08	53.06	51.96
0.2	7.98	8.92	7.02	0.99	53.01	51.92
0.3	9.26	9.43	7.01	1.03	53.02	51.95
Average	7.89	8.39	7.01	1.03	53.03	51.85

Table 2.4: Average computational time(seconds) of the 8-period numerical for different pivoting parameters

2.7.2 A 25-period test bed

As shown in Section 2.7.1 for the 8-period test bed, both the MP and the BS methods provide tight optimality gaps and acceptable computational efficiency. We now extend the 8-period test bed to 25 periods with larger instances. Demand of LCY1, LCY2, SIN1, SIN2, STA, and RAND are generated with expressions (2.20), (2.21), (2.22), (2.23), (2.24), and (2.25) in Fig. 2.5. Demand of EMP1, EMP2, EMP3 and EMP4 are derived from Strijbosch et al. (2011). Full details are given in Appendix 2.C. Assume that fixed ordering cost ranges in $\{500, 1000, 1500\}$, proportional ordering cost ranges in $\{0, 1\}$, penalty cost takes values $\{5, 10, 20\}$, and the coefficients of standard deviations are $\{0.1, 0.2, 0.3\}$.

$$d_t = \text{round}\left(\frac{190 \times e^{-(t-13)^2}}{2 \times 5^2}\right), \quad t = 1, 2, \dots, T \quad (2.20)$$

$$d_t = \text{round}\left(\frac{170 \times e^{-(t-13)^2}}{2 \times 6^2}\right), \quad t = 1, 2, \dots, T \quad (2.21)$$

$$d_t = \text{round}\left(70 \times \sin(0.8t) + 80\right), \quad t = 1, 2, \dots, T \quad (2.22)$$

$$d_t = \text{round}\left(30 \times \sin(0.8t) + 100\right), \quad t = 1, 2, \dots, T \quad (2.23)$$

$$d_t = 100, \quad t = 1, 2, \dots, T \quad (2.24)$$

$$d_t = \text{round}(\text{random}(0, 250)), \quad t = 1, 2, \dots, T \quad (2.25)$$

Figure 2.5: Expressions for generating demand data

We obtain optimal (s, S) parameters and record computational time obtained via the BS algorithm. For the first 15 periods we perform binary search with step size 1 in order to ensure fast convergence; for the last 10 periods, we adopt a step size 0.1 to enhance accuracy. The number of segments used in the piecewise linearisation is eleven. To estimate the cost of the policy obtained via our approximation, we simulate each instance ten thousand times in MATLAB.

We observe that the BS algorithm has high accuracy even for the large-size numerical experiments. We report detailed model accuracy in Table 2.9 in Appendix 2.D.

In Table 2.5, we summarise computational times of the BS model for different pivoting parameters. It is difficult to make a general remark with respect to demand patterns. An increase in fixed ordering cost significantly decreases the computational time. For instance, the computational time drops from 934.92 to 546.75 seconds as the fixed ordering cost increases from 500 to 1500. An increase

in proportional ordering cost decreases the computational time. In contrast, an increase in coefficient of variation increases the computational time. For instance, when the coefficient of variation rises from 0.1 to 0.2, the computational time increases from 679.34 to 809.34 seconds. On average, the computational time is 748.20 seconds and the standard deviation is 616.43 seconds.

Settings	Mean	standard deviation
Demand pattern		
LCY1	531.66	204.45
LCY2	740.73	322.92
SIN1	500.44	177.17
SIN2	1622.92	624.58
STA	1709.00	706.67
RAND	407.08	131.11
EMP1	633.09	126.63
EMP2	188.19	37.45
EMP3	974.93	305.16
EMP4	173.95	44.87
Fixed ordering cost		
500	934.92	811.90
1000	762.96	540.73
1500	546.75	341.41
Proportional ordering cost		
0	827.15	680.28
1	669.25	534.88
Penalty cost		
5	713.45	564.80
10	782.53	669.09
20	744.28	612.21
Coefficient of variation		
0.1	679.34	567.29
0.2	755.92	619.07
0.3	809.34	656.18
Average	748.20	616.43

Table 2.5: BS heuristic on a 25-period test bed, average computational time(seconds) with different parameter settings

2.8 Conclusion

In this paper we discussed two MINLP-based heuristics for tackling non-stationary stochastic lot-sizing problems under (s, S) policy.

Our first heuristic — the first MINLP heuristic for computing near-optimal

non-stationary (s, S) policy parameters — is based on mathematical programming models that can be solved by using off-the-shelf optimization packages. These MINLP models can be linearised via the approach discussed in Rossi et al. (2015) and can be implemented in OPL by adopting the `piecewise` expression.

Our second heuristic is a binary search strategy that leverages the aforementioned MINLP models and can tackle larger-size problems. However, this latter heuristic requires dedicated code.

We conducted an extensive computational study comprising 540 instances. We considered ten demand patterns, three fixed ordering costs, two proportional ordering cost, three penalty costs and three coefficients of variation.

We first conducted a numerical study on small instances (8-period). We investigated the performance of both models by contrasting costs of the policy obtained with our models against costs of the optimal policy obtained via the stochastic dynamic programming. Optimality gaps observed are generally below 0.3%. Our sensitivity analysis showed that the optimality gap is tighter when the demand keeps stable, and it deteriorates with the increase of the penalty cost and the coefficient of variation; both heuristics provide tighter gaps than those reported in the literature (Askin, 1981; Bollapragada and Morton, 1999).

The computational study carried out on larger instances (25-period) showed that the computational efficiency of the binary search approach is reasonable: around 748.20 seconds on average.

Appendix

2.A The piecewise OPL constraint

Rossi et al. (2015) piecewise linearised loss functions in constraints (2.9) and (2.10) by employing piecewise linear approximations based on Jensen’s and Edmundson-Madanski inequalities. An alternative strategy is to model these non-linear functions by exploring the `piecewise` syntax in OPL. By using this syntax, a piecewise function is specified by giving a set of slopes which represent the linear variation for each linear segment; a set of breakpoints at which slopes change; and the function value at a known point.

The `piecewise` syntax in OPL is given in Figure 2.6. W is the number of breakpoints of the piecewise function. `slope[i]` and `breakpoint[i]` denote slope and breakpoint of segment i . Segment i goes from breakpoint $(i-1)$ to breakpoint (i) . `<valuepoint>` is the function value at a known point `<knownpoint>`. Finally,

```

piecewise(i in 1..W){
slope[i] -> breakpoint[i];
slope[W+1]
}(<knownpoint>,<valuepoint>)<value>;

```

Figure 2.6: The syntax of the `piecewise` command in OPL

`<value>` represents the value at which we evaluate the function.

For the OPL `piecewise` syntax, there are three key components: slope, breakpoint, and function value at a known point. The following lemmas will demonstrate how to deduce their values. Let Ω be the support of ω . Let $(\Omega_i)_{i=1,\dots,W+1}$ be a partition of Ω in $W + 1$ segments.

Lemma 1. *The slope of i^{th} segment is written as*

$$l_i = \sum_{k=1}^{i-1} p_k, i \in \{1, 2, \dots, W + 1\},$$

where $p_i = Pr\{\omega \in \Omega_i\} = \int_{\Omega_i} g_\omega(t)dt$, $g_\omega(\cdot)$ denotes the probability density function of ω .

Proof. Observation from Rossi et al. (2014), Lemma 11. □

Lemma 2. *The i^{th} breakpoint can be written as*

$$X_i = E[\omega|\Omega_i], i \in \{1, 2, \dots, W\}.$$

Proof. Observation from Rossi et al. (2014), Lemma 11. □

Note that when ω follows a normal distribution with mean μ and standard deviation σ , then $\hat{L}_{up}(x, \omega) = \sigma \hat{L}_{up}(\frac{x-\mu}{\sigma}, Z)$, where Z follows a standard normal distribution, see Lemma 7 in Rossi et al. (2014).

Lemma 3. *Assume that the partition of Ω is symmetric with respect to 0, then the function value $\hat{L}_{up}(x, \omega)$ at point 0 can be written as follows.*

$$\hat{L}_{up}(0, \omega) = \begin{cases} -\sum_{k=1}^{\frac{W+1}{2}} p_k E[\omega|\Omega_k] + e_W, & W \text{ is odd} \\ -\frac{1}{2}(\sum_{k=1}^{\frac{W}{2}} p_k E[\omega|\Omega_k] + \sum_{k=1}^{\frac{W}{2}+1} p_k E[\omega|\Omega_k]) + e_W, & W \text{ is even} \end{cases}$$

where e_W represents the approximation error.

Proof. Since the partition of Ω is symmetric when W is odd, $x = 0$ is the central breakpoint. Hence, the function value at this breakpoint can be calculated directly. However, when W is even, the function value at point $x = 0$ is the average of nearest two symmetric breakpoints $X_{\frac{W}{2}}$ and $X_{\frac{W}{2}+1}$. \square

Following Lemma 1, 2 and 3, constraint (2.9) and (2.10) in Fig. 2.2 can be expressed as Eq. (2.26) and (2.27) in Fig. 2.7, for $t = 1, \dots, T$.

$$P_{jt} = 1 \rightarrow H_t = \text{piecewise}\{l_i \rightarrow X_i; 1\}(0, \hat{L}_{up}(0, d_{jt}))\tilde{I}_t, \\ i = 1, \dots, W; j = 1, \dots, t. \quad (2.26)$$

$$P_{jt} = 1 \rightarrow B_t = \text{piecewise}\{-1 + l_i \rightarrow X_i; 0\}(0, \hat{L}_{up}(0, d_{jt}))\tilde{I}_t, \\ i = 1, \dots, W; j = 1, \dots, t. \quad (2.27)$$

Figure 2.7: Rewriting holding and penalty costs by adopting `piecewise` syntax

2.B The MILP model

The joint MILP model to calculate near-optimal (s, S) policy parameters for the non-stationary stochastic lot-sizing problem is presented below.² In the joint MP model, constraints (2.30) represent the costs of controlling the system optimally when the initial inventory level is I_0^s ; constraints (2.41) denote the costs of controlling the system optimally when the initial inventory level is I_0^s , and no order is placed in period 1. These two constraints are connected via constraints (2.54) such that the order-up-to level S_1 and reorder point s_1 are approximated by I_0^S and I_0^s respectively.

²The loss function is piecewise linearized via constraints (2.37), (2.38), (2.49), and (2.50).

$$\min \left(\sum_{t=1}^T (K\delta_t^S + hH_t^S + bB_t^S) + c\tilde{I}_T^S + \sum_{t=2}^T (K\delta_t^s + hH_t^s + bB_t^s) + c\tilde{I}_T^s \right) \quad (2.28)$$

$$\text{Subject to, } t = 1, \dots, T \quad (2.29)$$

$$C_1^S(I_0^S) = \sum_{t=1}^T (K\delta_t^S + hH_t^S + bB_t^S) + c\tilde{I}_T^S + c \sum_{t=1}^T \tilde{d}_t - cI_0^S \quad (2.30)$$

$$\tilde{I}_t^S + \tilde{d}_t - \tilde{I}_{t-1}^S \geq 0 \quad (2.31)$$

$$\delta_t^S = 0 \rightarrow \tilde{I}_t^S + \tilde{d}_t - \tilde{I}_{t-1}^S = 0 \quad (2.32)$$

$$\sum_{j=1}^t P_{jt}^S = 1 \quad (2.33)$$

$$P_{jt}^S \geq \delta_j^S - \sum_{k=j+1}^t \delta_k, \quad j = 1, \dots, t \quad (2.34)$$

$$\delta_1^S = 1 \quad (2.35)$$

$$I_0^S = \tilde{I}_1^S + \tilde{d}_1 \quad (2.36)$$

$$H_t^S \geq (I_t^S + \sum_{j=1}^t d_{jt} P_{jt}^S) \sum_{k=1}^i p_k - \sum_{j=1}^t \left(\sum_{k=1}^i p_k E[d_{jt}|\Omega_i] - e_W \right) P_{jt}^S, \quad i = 1, \dots, W \quad (2.37)$$

$$B_t^S \geq -I_t^S + (I_t^S + \sum_{j=1}^t d_{jt} P_{jt}^S) \sum_{k=1}^i p_k - \sum_{j=1}^t \left(\sum_{k=1}^i p_k E[d_{jt}|\Omega_i] - e_W \right) P_{jt}^S, \quad i = 1, \dots, W \quad (2.38)$$

$$P_{jt}^S \in \{0, 1\}, \quad j = 1, \dots, t \quad (2.39)$$

$$\delta_t^S \in \{0, 1\} \quad (2.40)$$

$$G_1^S(I_0^S) = (hH_1^S + bB_1^S) + \sum_{t=2}^T (K\delta_t^S + hH_t^S + bB_t^S) + c\tilde{I}_T^S + c \sum_{t=1}^T \tilde{d}_t - cI_0^S \quad (2.41)$$

$$\tilde{I}_t^S + \tilde{d}_t - \tilde{I}_{t-1}^S \geq 0 \quad (2.42)$$

$$\delta_t^S = 0 \rightarrow \tilde{I}_t^S + \tilde{d}_t - \tilde{I}_{t-1}^S = 0 \quad (2.43)$$

$$\sum_{j=1}^t P_{jt}^S = 1 \quad (2.44)$$

$$P_{jt}^S \geq \delta_j^S - \sum_{k=j+1}^t \delta_k^S, \quad j = 1, \dots, t \quad (2.45)$$

$$\delta_1^S = 0 \quad (2.46)$$

$$P_{jt}^S = 1 \rightarrow H_t^S = \text{piecewise}\{l_i \rightarrow X_i; 1\}(0, \hat{L}_{up}(0, d_{jt}))\tilde{I}_t^S, \quad \begin{matrix} i = 1, \dots, W \\ j = 1, \dots, t \end{matrix} \quad (2.47)$$

$$P_{jt}^S = 1 \rightarrow B_t^S = \text{piecewise}\{-1 + l_i \rightarrow X_i; 0\}(0, \hat{L}_{up}(0, d_{jt}))\tilde{I}_t^S, \quad \begin{matrix} i = 1, \dots, W \\ j = 1, \dots, t \end{matrix} \quad (2.48)$$

$$H_t^S \geq (I_t^S + \sum_{j=1}^t d_{jt} P_{jt}^S) \sum_{k=1}^i p_k - \sum_{j=1}^t \left(\sum_{k=1}^i p_k E[d_{jt}|\Omega_i] - e_W \right) P_{jt}^S, \quad i = 1, \dots, W \quad (2.49)$$

$$B_t^S \geq -I_t^S + (I_t^S + \sum_{j=1}^t d_{jt} P_{jt}^S) \sum_{k=1}^i p_k - \sum_{j=1}^t \left(\sum_{k=1}^i p_k E[d_{jt}|\Omega_i] - e_W \right) P_{jt}^S, \quad i = 1, \dots, W \quad (2.50)$$

$$P_{jt}^S \in \{0, 1\}, \quad j = 1, \dots, t \quad (2.51)$$

$$\delta_t^S \in \{0, 1\} \quad (2.52)$$

$$I_0^S \leq I_0^s \quad (2.53)$$

$$G_1^S(I_0^S) = C_1^S(I_0^S) + c \cdot (I_0^S - I_0^s) \quad (2.54)$$

2.C Test bed

Expected demand patterns under the eight period computational study are displayed in Table 2.6. The demand of each period under the twenty-five periods numerical example is shown in Table 2.7. The first column represents period

indexes; the rest columns denote various demands.

Period	LCY1	LCY2	SIN1	SIN2	STA	RAND	EMP1	EMP2	EMP3	EMP4
1	15	3	15	12	10	2	5	4	11	18
2	16	6	4	7	10	4	15	23	14	6
3	15	7	4	7	10	7	26	28	7	22
4	14	11	10	10	10	3	44	50	11	22
5	11	14	18	13	10	10	24	39	16	51
6	7	15	4	7	10	10	15	26	31	54
7	6	16	4	7	10	3	22	19	11	22
8	3	15	10	12	10	3	10	32	48	21

Table 2.6: Demand data of the 8-period computational analysis

Period	LCY1	LCY2	SIN1	SIN2	STA	RAND	EMP1	EMP2	EMP3	EMP4
1	11	23	130	122	100	178	2	47	44	49
2	17	32	150	130	100	178	51	81	116	188
3	26	42	127	120	100	136	152	236	264	64
4	38	55	76	98	100	211	467	394	144	279
5	53	70	27	77	100	119	268	164	146	453
6	71	86	10	70	100	165	489	287	198	224
7	92	103	36	81	100	47	446	508	74	223
8	115	120	88	103	100	100	248	391	183	517
9	138	136	136	124	100	62	281	754	204	291
10	159	150	149	130	100	31	363	694	114	547
11	175	161	121	118	100	43	155	261	165	646
12	186	168	68	95	100	199	293	195	318	224
13	190	170	22	75	100	172	220	320	119	215
14	186	168	11	71	100	96	93	111	482	440
15	175	161	42	84	100	69	107	191	534	116
16	159	150	96	107	100	8	234	160	136	185
17	138	136	140	126	100	29	124	55	260	211
18	115	120	148	129	100	135	184	84	299	26
19	92	103	114	115	100	97	223	58	76	55
20	71	86	60	91	100	70	101	0	218	0
21	53	70	18	73	100	248	123	0	323	0
22	38	55	14	72	100	57	99	0	102	0
23	26	42	50	87	100	11	31	0	174	0
24	17	32	104	110	100	94	82	0	284	0
25	11	23	144	127	100	13	0	0	0	0

Table 2.7: Demand data of the 25-period computational analysis

2.D Model accuracy

We employ the index of model accuracy ($= \frac{|\text{model result} - \text{simulation result}|}{\text{simulation result}} \times 100\%$) to evaluate the cost measure. We report the model accuracy of the 8-period

numerical experiment in Table. 2.8, and the 25-period numerical experiment in Table. 2.9.

Modelling methods	MP	BS
Demand pattern		
LCY1	1.52	0.66
LCY2	7.47	3.42
SIN1	0.99	0.37
SIN2	0.84	0.30
STA	1.25	0.66
RAND	4.57	2.10
EMP1	8.75	4.50
EMP2	6.82	3.05
EMP3	1.83	0.81
EMP4	2.59	0.73
Fixed ordering cost		
200	3.14	1.36
300	3.71	1.66
400	4.15	1.96
Proportional ordering cost		
0	4.00	0.58
1	3.33	4.72
Penalty cost		
5	5.29	2.47
10	3.27	1.33
20	2.42	1.18
Coefficient of variation		
0.1	2.94	1.33
0.2	3.74	1.60
0.3	4.31	2.05
Average gap	3.66	1.66

Table 2.8: Model accuracy of the 8-period numerical experiments

Modelling method	BS
Demand pattern	
LCY1	2.32
LCY2	2.97
SIN1	2.65
SIN2	2.50
STA	1.90
RAND	2.81
EMP1	4.15
EMP2	5.19
EMP3	3.79
EMP4	5.55
Fixed ordering cost	
500	3.27
1000	3.46
1500	3.42
Proportional ordering cost	
0	3.52
1	3.24
Penalty cost	
5	2.56
10	3.23
20	4.34
Coefficient of variation	
0.1	1.68
0.2	3.13
0.3	5.34
Average gap	3.38

Table 2.9: Accuracy of the 25-period numerical experiments

Chapter 3

Paper II: (R, S) policy with correlated demand

Abstract

This paper considers the single-item single-stock location non-stationary stochastic lot-sizing problem under correlated demand. By operating under a nonstationary (R, S) policy, in which R denotes the length of the replenishment interval and S the associated order-up-to level, we introduce a mixed integer linear programming (MILP) model which can be easily implemented by using off-the-shelf optimisation software. Our modelling strategy can tackle a wide range of time-series-based demand processes, such as autoregressive (AR), moving average (MA), autoregressive moving average (ARMA), and autoregressive with autoregressive conditional heteroskedasticity process (AR-ARCH). In an extensive computational study, we compare the performance of our model against the optimal policy obtained via stochastic dynamic programming. Our results demonstrate that the optimality gap of our approach averages 2.28% and that computational performance is good.

3.1 Introduction

Stochastic lot sizing is an important area of research in inventory control (Graves, 1999). Scarf's pioneering work (Scarf, 1960) proved the optimality of (s, S) policies for a class of dynamic inventory models. Since then a sizeable literature fo-

cused on the computation of the optimal policy parameters (see e.g. (Veinott Jr, 1965; Askin, 1981; Federgruen and Zipkin, 1984)). The (s, S) policy allows decision makers to decide dynamically at each time period whether or not to place an order, by checking if the inventory level is below the reorder threshold s ; and how much to order, by “topping” inventory up to level S . However, as pointed out in Tarim and Smith (2008), this policy performs poorly in terms of “nervousness” of the control action, i.e. it suffers from planning instability. In this regard, Bookbinder and Tan (1988) discussed the other two policies: static, and static-dynamic uncertainty. The static uncertainty strategy, also known as (R, Q) , enables decision makers to decide the timing (R) and size (Q) of replenishments at the beginning of the planning horizon. The static-dynamic uncertainty strategy, known as (R, S) policy, provides an effective means for reducing 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 this paper, we focus our attention on the (R, S) policy.

Several approaches for computing optimal (R, S) policy parameters have been proposed, e.g. (Bookbinder and Tan, 1988; Tarim and Kingsman, 2004, 2006; Rossi et al., 2015). A common assumption in all these studies is that random demand in each period is independent of demand in other periods. However, as discussed in Song and Zipkin (1993), environmental factors, such as economic conditions, market conditions, and any exogenous conditions, have major effects on the demand for a product, the supply, and the cost structure. In this regards, the goal of this paper is to *relax the assumption of independence of demand in different periods*.

Correlated demand has been previously investigated in the inventory literature. Authors attempted to either prove the optimality of (s, S) policy, or compute optimal policy parameters with different types of demand correlations over the planning horizon. However, to the best of our knowledge, no study on computing (R, S) policies under time-series-based demand processes exists.

In this paper, we consider a periodic-review single-item single-stocking location lot-sizing problem under *non-stationary stochastic correlated demand*. We build upon Rossi et al. (2015), which discussed a mixed integer linear programming (MILP) heuristic for approximating the optimal (R, S) policies under stochastic demand independent from period to period. We leverage properties of conditional distributions, and present an MILP-based heuristic for approximating

optimal (R, S) policies under normally distributed demand featuring correlation across periods as well as under a collection of time-series-based demand processes. Our approach offers a stable replenishment plan while effectively hedging against uncertainty. Our model can be easily implemented and solved by using off-the-shelf mathematical programming packages such as IBM ILOG optimisation studio.

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

- We developed a stochastic model which captures the (R, S) policy under correlated demand — to the best of our knowledge this is the first time the (R, S) policy has been expressed in the form of a functional equation.
- We present an MILP-based heuristic for approximating optimal (R, S) policies under normally distributed demand featuring correlation across periods; our MILP model can be easily solved by using off-the-shelf software.
- We illustrate how to adapt the model to a collection of time-series-based demand processes: the autoregressive (AR) process, the moving-average (MA) process, the autoregressive moving-average (ARMA) process, and the autoregressive with autoregressive conditional heteroskedasticity (AR-ARCH) process.
- Our computational experiments demonstrate that the MILP heuristic provides tight optimality gaps and good computational times.

The rest of this paper is organised as follows. Section 3.2 surveys the related literature. Section 3.3 discusses relevant properties of multivariate normal distribution, and the stochastic dynamic programming (SDP) formulation under correlated demand. Section 3.4 derives a stochastic model which captures the (R, S) policy. Section 3.5 presents our MILP model under correlated demand. Section 3.6 shows how our MILP model can be extended to cover a collection of time-series-based demand processes. In Section 3.7 we present an extensive computational study. Finally, we draw conclusions in Section 3.8.

3.2 Literature review

In this section, we first survey literature on the (R, S) policy addressing the case of independently and identically distributed demand in each period. We then survey literature on correlated demand; in this stream of literature most studies focused on establishing the optimality of (s, S) policies under a range of time-series-based demand processes. This paper differs from existing research by considering (R, S)

policy under normally distributed demand featuring correlation across periods, as well as under a collection of time-series-based demand processes.

The (R, S) policy with independently and identically distributed demand has been extensively studied. In Bookbinder and Tan (1988) the authors proposed a two-stage deterministic equivalent heuristic which first fixes replenishment periods, and then determines order quantities for a single item inventory system with fixed and proportional ordering costs, holding cost, and service level constraints. Later, Tarim and Kingsman (2004) formulated a mixed integer programming (MIP) model for determining both timing and quantity of orders simultaneously. In a follow-up study, Tarim and Kingsman (2006) incorporated penalty costs. Tarim et al. (2011) relaxed the original MIP model of Tarim and Kingsman (2004), and solved it as a shortest path problem which does not require the use of any MIP or constraint programming (CP) commercial solver. In addition, Özen et al. (2012) showed a DP-based algorithm for solving small-size problems, and an approximation heuristic and a relaxation heuristic for tackling larger-size problems; Tunc et al. (2014) suggested a deterministic equivalent MIP model. Recently, Rossi et al. (2015) generalised the discussions above, developed a unified MILP model with service level constraints, penalty costs, and lost sale settings by adopting the piecewise linear approximation technique in Rossi et al. (2014). Although various efficient modelling approaches were proposed, they generally assume that demand is independently and identically distributed, which is often unrealistic. In this paper we build upon Rossi et al. (2015) and present an MILP-based heuristic for approximating (R, S) policies under correlated demand.

Literature on correlated demand can be roughly classified into two streams. The first stream focused on establishing the optimality of (s, S) policy; while the second stream focused on performances of different policies with different time-series-based demand processes.

In relation to establishing the optimality of (s, S) policy under correlated demand, Iglehart (1962) studied the case of Markovian demand considering fixed and unit ordering cost, holding cost and shortage cost. Sethi and Cheng (1997) established optimality of (s, S) policy for a generalization of classical inventory models, including non-ordering periods, finite storage capacities, and service levels. Beyer and Sethi (1997) incorporated convex surplus cost into the model and proved the optimality from the viewpoint of minimising the long-run average cost of inventory/backlog and ordering. Lian et al. (2009) proposed the first perishable inventory model with Markovian renewal demand, and proved the optimal policy is (s, S) type. Multiechelon models incorporating Markov-modulated demand are

discussed in (Chen and Song, 2001; Hu et al., 2016). Other studies have shown that the optimality of the (s, S) policy can be generalized to cases involving unbounded Markovian demand (Beyer et al., 1998), unreliable suppliers (Özekici and Parlar, 1999), and polynomial growth demand, returns, and cost functions (Li, 2013).

Regarding performances of different policies under time-series-based demand processes, a widely adopted policy is the “base stock” policy. Under this policy, if the opening inventory level is less than the base stock level, then an order is issued to increase its inventory level to the base stock level; otherwise, no order is issued. Johnson and Thompson (1975) proved the optimality of the base stock policy for single-item periodic ordering systems with proportional holding and stock-out costs and zero lead time for both Autoregressive (AR) and Moving Average (MA) demand processes under the condition that demand falls between certain lower and upper bounds, but without actually computing the optimal values.

Graves (1999) developed a single-item inventory model under a deterministic lead-time and an integrated moving average process. Dong and Lee (2003) approximated the optimal base stock level when the demand is time-correlated with a Martingale model of forecast evolution, and provided a simple, easy-to-compute closed form expression for base stock level and average system costs under, in particular, the AR(1) process.

Further policies under time-series-based demand processes are the following. Ray (1981) focused on calculating the “reorder level” policy with random lead time, and AR and MA demand processes. Fotopoulos et al. (1988) presented a straightforward method for computing optimal policies with correlated AR and MA demand process, and arbitrary lead times for the (s, S) policy. Recently, Carrizosa et al. (2016) adopted a robust approach to explore the single-item news-vendor problem with AR(P) demand processes. A close-form expression for computing optimal order quantity is found for the AR(1) process; for the remaining higher order AR processes, the problem is expressed as a solvable non-linear convex optimization program. On the basis of our survey, no study has been found in the literature that addresses the (R, S) policy with time-series-based demand processes.

Capturing the behaviour of the demand process is integral to the analysis of inventory management systems (Nasr and Maddah, 2015). All studies we surveyed either address (R, S) policy with independently and identically distributed demand, or investigated specific demand correlations under the (s, S) policy, the

base stock policy, or the reorder level policy. The contribution of this paper is to present an MILP-based heuristic for approximating the optimal (R, S) policies with a collection of time-series-based demand processed, which has not been addressed yet in the literature.

3.3 A stochastic dynamic programming formulation

We consider a stochastic lot-sizing problem over a T -period planning horizon. Demand d_t in each period $t = 1, \dots, T$ is a normally distributed random variable with probability density function $g_{d_t}(\cdot)$. We assume that distributions of demand in successive periods are not identically distributed, and generally are correlated. A full list of symbols is available in Appendix 3.A.

Let d be a n -variate multivariate normal random variable with mean μ and variance-covariance matrix Σ , abbreviated $d \sim \mathcal{MVN}(\mu, \Sigma)$, where

$$d = \begin{bmatrix} d_1 \\ d_2 \\ \vdots \\ d_n \end{bmatrix} \quad \mu = \begin{bmatrix} \mu_1 \\ \mu_2 \\ \vdots \\ \mu_n \end{bmatrix} \quad \Sigma = \begin{bmatrix} \text{Var}(d_1) & \text{Cov}(d_1, d_2) & \dots & \text{Cov}(d_1, d_n) \\ \text{Cov}(d_2, d_1) & \text{Var}(d_2) & \dots & \text{Cov}(d_2, d_n) \\ \vdots & \vdots & \ddots & \vdots \\ \text{Cov}(d_n, d_1) & \text{Cov}(d_n, d_2) & \dots & \text{Var}(d_n) \end{bmatrix}.$$

We present two fundamental theorems for conditional distribution and linear transformation.

Theorem 3.3.1 (Conditional distribution (Billingsley, 2008)). *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 $\mu = [\mu_1 \ \mu_2]^T$ and variance-covariance matrix $\Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}$. Then the conditional distribution of d_2 given $d_1 = \eta_1$ is multivariate normal with*

$$E[d_2 | d_1 = \eta_1] = \mu_2 + \Sigma_{21} \Sigma_{11}^{-1} (\eta_1 - \mu_1) \quad (3.1)$$

$$\text{Cov}(d_2 | d_1 = \eta_1) = \Sigma_{22} - \Sigma_{21} \Sigma_{11}^{-1} \Sigma_{12}. \quad (3.2)$$

Theorem 3.3.2 (Linear transformations (Taboga, 2012)). *If $d \sim \mathcal{MVN}(\mu, \Sigma)$, then any linear combinations of the d_i for $i \in \{1, \dots, n\}$, say $a^T d = a_1 d_1 + \dots + a_n d_n$, is normally distributed as $a^T d \sim \mathcal{MVN}(a^T \mu, a^T \Sigma a)$.*

Let $\mathcal{D}_t = (\eta_1, \dots, \eta_{t-1})$ represent demand realisations at the beginning of

period t , then the demand of period t has the conditional probability density function $g_{d_t}(\zeta_t|\mathcal{D}_t)$.

In what follows, we define variables I_t , and Q_t . I_t , the inventory level at the end of period t , and the opening inventory level of period $t + 1$ before replenishment. Q_t , the ordering quantity at the beginning of period t . Let I_0 represent the given initial inventory level at the beginning of the planning horizon.

We further assume that orders are placed at the beginning of each time period, and delivered instantaneously. There exist ordering costs $c(\cdot)$ comprising a fixed ordering cost K for placing an order, and a linear ordering cost c proportional to the order quantity Q_t ; which takes the following form.

$$c(Q_t) = \begin{cases} K + c \cdot Q_t & Q_t > 0 \\ 0 & Q_t = 0 \end{cases} \quad (3.3)$$

Additionally, 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 unit of unmet demand.

Given the above problem description, the objective is to schedule ordering plans so as to minimize the expected total cost. The problem can be formulated as a stochastic dynamic program (Bellman, 1966) comprising the following elements.

1. **Stage.** A stage represents a time period $t = \{1, \dots, T\}$ for a T -period stochastic lot-sizing problem.
2. **State.** Let \mathcal{S}_t denote the state of the system at the beginning of period t before replenishment. State $\mathcal{S}_t = \{I_{t-1}, \mathcal{D}_t\}$ includes the opening inventory level I_{t-1} of period t , and the realised demand information set $\mathcal{D}_t = (\eta_1, \dots, \eta_{t-1})$.
3. **Action.** An action means to schedule an order with quantity Q_t at the beginning of period t , $Q_t \in [0, \infty)$.
4. **Immediate cost.** Let $f_t(I_{t-1}, \mathcal{D}_t, Q_t)$ denote the expected immediate cost comprising ordering, holding, and penalty costs, given state $\mathcal{S}_t = \{I_{t-1}, \mathcal{D}_t\}$.

$$\begin{aligned} f_t(I_{t-1}, \mathcal{D}_t, Q_t) &= c(Q_t) + h \int_{d_t} \max(I_{t-1} + Q_t - \zeta_t, 0) g_{d_t}(\zeta_t|\mathcal{D}_t) d(\zeta_t) \\ &\quad + b \int_{d_t} \max(\zeta_t - I_{t-1} - Q_t, 0) g_{d_t}(\zeta_t|\mathcal{D}_t) d(\zeta_t) \\ &= c(Q_t) + h E_{d_t}[\max(I_{t-1} + Q_t - d_t, 0) | \mathcal{D}_t] \\ &\quad + b E_{d_t}[\max(d_t - I_{t-1} - Q_t, 0) | \mathcal{D}_t] \end{aligned} \quad (3.4)$$

5. **Objective function.** Let $C_t(I_{t-1}, \mathcal{D}_t)$ denote the expected total cost of an optimal policy over period t, \dots, T with state $\mathcal{S}_t = \{I_{t-1}, \mathcal{D}_t\}$. Then $C_t(I_{t-1}, \mathcal{D}_t)$ can be written as, for $t = 1, \dots, T - 1$,

$$\begin{aligned} C_t(I_{t-1}, \mathcal{D}_t) &= \min_{Q_t \geq 0} \left\{ f_t(I_{t-1}, \mathcal{D}_t, Q_t) + \int_{d_t} C_{t+1}(I_{t-1} + Q_t - \zeta_t, \mathcal{D}_{t+1}) g_{d_t}(\zeta_t | \mathcal{D}_t) d(\zeta_t) \right\} \\ &= \min_{Q_t \geq 0} \left\{ f_t(I_{t-1}, \mathcal{D}_t, Q_t) + E_{d_t}[C_{t+1}(I_{t-1} + Q_t - d_t, \mathcal{D}_{t+1}) | \mathcal{D}_t] \right\}, \end{aligned} \quad (3.5)$$

where

$$C_T(I_{T-1}, \mathcal{D}_T) = \min_{Q_T \geq 0} \{ f_T(I_{T-1}, \mathcal{D}_T, Q_T) \} \quad (3.6)$$

represents the boundary condition.

Example. We now introduce a 4-period example. Demand d_t in successive periods are correlated with covariance coefficient $\rho = 0.5$, d_t in each time period are normally distributed with means $\mu_t = 20, 40, 60, 40$, and standard deviations $0.25\mu_t$. Other parameters are $K = 100$, $h = 1$, $b = 10$, $c = 0$, and $I_0 = 0$. By solving the stochastic dynamic program, we obtain an expected total cost equal to 378.06.

3.4 Towards an (R, S) policy

The (R, S) policy, proposed by Bookbinder and Tan (1988), features two parameters: R and S . Under this policy, the review times R and the respective order-up-to levels S are fixed at the beginning of the planning horizon. However, actual ordering quantities are decided at the beginning of each review period to reach the order-up-to level.

In this section we introduce a stochastic model which captures the (R, S) policy. We begin by reformulating the stochastic dynamic program with fixed timing of replenishments, defined as the “static-dynamic uncertainty” policy (Bookbinder and Tan, 1988). We then fix the order-up-to level of replenishments, thus obtaining a stochastic dynamic program under (R, S) policy (Tarim and Kingsman, 2004). Finally, we produce a stochastic model which captures the (R, S) policy.

3.4.1 “Static-dynamic uncertainty” policy

The “static-dynamic uncertainty” policy (Bookbinder and Tan, 1988) requires the timing of replenishments to be fixed at the beginning of the planning horizon, while the actual replenishment quantities are decided at the beginning of each ordering period.

We first introduce a binary variable δ_t , for $t = 1, \dots, T$, which takes value 1 if a replenishment is placed in period t and 0 otherwise. Then, the ordering cost in Eq. (3.3) is replaced as follows,

$$c(Q_t) = K\delta_t + c \cdot Q_t, \quad (3.7)$$

for $t = 1, \dots, T$, and $Q_t \geq 0$. Thus, the timing of replenishments are given by the values of t such that $\delta_t = 1$; the quantities of replenishments are given by values of Q_t .

At the beginning of the planning horizon, before demand information becomes available, the system state is $\mathcal{S}_1 = \{I_0\}$. The objective is to decide the ordering periods over the planning horizon, and the ordering quantity of period 1 so as to minimise the expected total cost. Let $\hat{C}_1(I_0)$ represent the expected total cost of an optimal policy over periods $1, \dots, T$, given the initial inventory level I_0 at the beginning of period 1.

$$\hat{C}_1(I_0) = \min_{\substack{\delta_1, \dots, \delta_T \\ 0 \leq Q_1 \leq M\delta_1}} \left\{ f_1(I_0, Q_1) + \int_{d_1} \hat{C}_2(I_0 + Q_1 - \zeta_t, \mathcal{D}_2) g_{d_1}(\zeta_t) d(\zeta_t) \right\}. \quad (3.8)$$

Note that the constraint $0 \leq Q_1 \leq M\delta_1$ represents the order quantity must lie between 0 and a sufficiently large number, M . If an order is placed in period 1, i.e. $\delta_1 = 1$, the order quantity must be a non-negative integer; otherwise, it is 0.

Since the timing of replenishments are decided at the beginning of period 1, the objective in period t , for $t = 2, \dots, T$, is to decide replenishment quantities such that the expected total future cost is minimised under the given system state $\mathcal{S}_t = \{I_{t-1}, \mathcal{D}_t\}$. Therefore, the expected total cost $\hat{C}_t(I_{t-1}, \mathcal{D}_t)$ is, for $t = 2, \dots, T - 1$

$$\hat{C}_t(I_{t-1}, \mathcal{D}_t) = \min_{0 \leq Q_t \leq M\delta_t} \left\{ f_t(I_{t-1}, \mathcal{D}_t, Q_t) + \int_{d_t} \hat{C}_{t+1}(I_{t-1} + Q_t - \zeta_t, \mathcal{D}_{t+1}) g_{d_t}(\zeta_t | \mathcal{D}_t) d(\zeta_t) \right\}, \quad (3.9)$$

where

$$\hat{C}_T(I_{T-1}, \mathcal{D}_T) = \min_{0 \leq Q_T \leq M\delta_T} \{f_T(I_{T-1}, \mathcal{D}_T, Q_T)\} \quad (3.10)$$

represents the boundary condition.

3.4.2 Modelling the (R, S) policy via stochastic dynamic programming

In the last section, under the “static-dynamic uncertainty” policy (Bookbinder and Tan, 1988), the replenishment periods are decided at the beginning of the planning horizon, and the actual order quantities are decided at the beginning of each ordering period. In this section, we follow the (R, S) policy introduced in Tarim and Kingsman (2004), where not only the timing of replenishments (R), but also the corresponding order-up-to levels (S) are fixed at the beginning of the planning horizon. Thus, the actual ordering quantities Q_t are uniquely decided, at the beginning of each ordering period, by the associated order-up-to levels S_t and opening inventory levels I_{t-1} , i.e. $Q_t = S_t - I_{t-1}$ if $\delta_t = 1$, and $Q_t = 0$ otherwise.¹

At the beginning of the planning horizon, before demand information becomes available, the system state is $\mathcal{S}_t = \{I_0\}$. The objective is to determine $\{\delta_1, \dots, \delta_T\}$ and $\{S_1, \dots, S_T\}$ so as to minimise the expected total cost. Let $\check{C}_1(I_0)$ represent the expected total cost of an optimal policy over periods $1, \dots, T$, given the initial inventory level I_0 at the beginning of period 1.

$$\check{C}_1(I_0) = \min_{\delta_1, \dots, \delta_T} \min_{\substack{S_1, \dots, S_T \\ Q_1 = (S_1 - I_0)\delta_1}} \{f_1(I_0, Q_1) + \int_{d_1} \check{C}_2(I_0 + Q_1 - \zeta_t, \mathcal{D}_2) g_{d_1}(\zeta_t) d(\zeta_t)\} \quad (3.11)$$

Note that the actual order quantity of period 1 is uniquely decided by $Q_1 = S_1 - I_0$ if a replenishment is placed, and 0 otherwise.

Since replenishment review periods and order-up-to levels are decided at the beginning of period 1, the objective in period t , for $t = 2, \dots, T$, is to determine the ordering quantity under the ordering schedule fixed at the beginning of the planning horizon. Therefore, the expected total cost over period t, \dots, T with given opening inventory level I_{t-1} and realised demand information set \mathcal{D}_t is as

¹Note that the probability that the opening inventory level I_{t-1} is greater than the order-up-to level S_t is generally small and can be safely neglected Rossi et al. (2008).

follows, for $t = 2, \dots, T - 1$,

$$\begin{aligned} \check{C}_t(I_{t-1}, \mathcal{D}_t) = & \min_{Q_t = (S_t - I_{t-1})\delta_t} \{f_t(I_{t-1}, \mathcal{D}_t, Q_t) \\ & + \int_{d_t} \check{C}_{t+1}(I_{t-1} + Q_t - \zeta_t, \mathcal{D}_{t+1})g_{d_t}(\zeta_t|\mathcal{D}_t)d(\zeta_t)\}, \end{aligned} \quad (3.12)$$

where

$$\check{C}_T(I_{T-1}, \mathcal{D}_T) = \min_{Q_T = (S_T - I_{T-1})\delta_T} \{f_T(I_{T-1}, \mathcal{D}_T, Q_T)\} \quad (3.13)$$

represents the boundary condition.

Since the ordering schedule is decided at the beginning of the planning horizon, the function $\check{C}_t(I_{t-1}, \mathcal{D}_t)$, for $t = 2, \dots, T$, only represents its linear relationship with order quantity Q_t . Then, the ‘‘min’’ symbol in Eq. (3.12)-(3.13) can be dropped. Therefore, the stochastic dynamic program under (R, S) policy can be rewritten as follows.

$$\check{C}_1(I_0) = \min_{\delta_1, \dots, \delta_T} \min_{\substack{S_1, \dots, S_T \\ Q_t = (S_t - I_{t-1})\delta_t}} \{f_1(I_0, Q_1) + \int_{d_1} \check{C}_2(I_0 + Q_1 - \zeta_t, \mathcal{D}_2)g_{d_1}(\zeta_t)d(\zeta_t)\} \quad (3.14)$$

$$\check{C}_t(I_{t-1}, \mathcal{D}_t) = f_t(I_{t-1}, \mathcal{D}_t, Q_t) + \int_{d_t} \check{C}_{t+1}(I_{t-1} + Q_t - \zeta_t, \mathcal{D}_{t+1})g_{d_t}(\zeta_t|\mathcal{D}_t)d(\zeta_t) \quad (3.15)$$

where $t = 2, \dots, T - 1$, and

$$\check{C}_T(I_{T-1}, \mathcal{D}_T) = f_T(I_{T-1}, \mathcal{D}_T, Q_T) \quad (3.16)$$

represents the boundary condition.

3.4.3 A stochastic programming model under (R, S) policy

In this section we reformulate the stochastic dynamic program as a stochastic programming model capturing the (R, S) policy. This reformulation is done by compacting Eq. (3.14)-(3.16), and replacing the immediate costs with Eq. (3.4). Then the stochastic programming formulation is given in Fig. 3.1.

The objective (3.17) is to decide the timing and order-up-to level of replenishment at the beginning of the planning horizon so as to minimise the expected

$$\begin{aligned} \check{C}_1(I_0) = & \min_{\delta_1, \dots, \delta_T} \min_{S_1, \dots, S_T} \int_{d_1} \cdots \int_{d_T} \sum_{t=1}^T \left(K\delta_t + c \cdot Q_t + h \max(I_{t-1} + Q_t - \zeta_t, 0) \right. \\ & \left. + b \max(\zeta_t - I_{t-1} - Q_t, 0) \right) g_{d_1}(\zeta_1 | \mathcal{D}_1) \cdots g_{d_T}(\zeta_T | \mathcal{D}_T) d(\zeta_1) \cdots d(\zeta_T) \end{aligned} \quad (3.17)$$

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

$$Q_t = (S_t - I_{t-1})\delta_t \quad (3.18)$$

$$I_t = I_0 + \sum_{i=1}^t (Q_i - \zeta_i) \quad (3.19)$$

$$Q_t, S_t \geq 0, I_t \in \mathcal{R}, \delta_t \in \{0, 1\} \quad (3.20)$$

Figure 3.1: A stochastic program under (R, S) policy with correlated demand

total cost comprising ordering, holding, and penalty costs. Constraints (3.18) ensure the ordering quantity must be equal to order-up-to level S_t , minus the opening inventory level I_{t-1} if an order is placed, and 0 otherwise. Constraints (3.19) are the inventory conservation constraints, the closing inventory I_t must be equal to the initial inventory level, plus all orders received, minus all conditional demand realised up to period t . Constraints (3.20) set order quantity and order-up-to level to be non-negative; inventory level could be any real number; δ_t is a binary variable.

We now simplify the stochastic program in Fig. 3.1 as follows by applying the law of total expectation (Weiss, 2006).

$$\begin{aligned} \check{C}_1(I_0) = & \min_{\delta_1, \dots, \delta_T} \int_{d_1} \cdots \int_{d_T} \sum_{t=1}^T \left(K\delta_t + c \cdot Q_t + h \max(I_{t-1} + Q_t - \zeta_t, 0) \right. \\ & \left. + b \max(\zeta_t - I_{t-1} - Q_t, 0) \right) g_{d_1}(\zeta_1) \cdots g_{d_T}(\zeta_T) d(\zeta_1) \cdots d(\zeta_T) \end{aligned} \quad (3.21)$$

The correlated demand case is similar to the independent demand case, which has been extensively discussed in Rossi et al. (2015). We will demonstrate how to approximate it in the next section.

3.5 Towards a Mixed Integer Linear Programming model

In this section we present a Mixed Integer Linear Programming (MILP) model for computing (R, S) policies. We begin by illustrating in Section 3.5.1 how to model a single replenishment cycle over periods i, \dots, j as well as multiple replenishment cycles. We then present in Section 3.5.2 an MILP model for computing optimal (R, S) policy parameters under correlated demand.

3.5.1 Fixed replenishment cycle problem

Consider a single replenishment cycle over periods i, \dots, j , where the only replenishment is placed at the beginning of period i with order-up-to level S_i , and the initial inventory level is I_{i-1} .

Let $d = [d_i \dots d_j] \in \mathcal{R}^{j-i+1}$ be a random vector. We assume that d has a multivariate normal distribution with mean $\tilde{d} = E[d] = [\tilde{d}_i \dots \tilde{d}_j]$, and variance-covariance matrix $\Sigma = \text{Cov}(d_m, d_n) = E[(d_m - \tilde{d}_m)(d_n - \tilde{d}_n)]$, $m = i, \dots, j$, and $n = i, \dots, j$.

Let a random variable d_{it} represent the convolution $d_i + \dots + d_t$, for $t = i, \dots, j$. Since the vector $[d_i \dots d_t]$ has multivariate normal distribution, d_{it} is normally distributed with mean \tilde{d}_{it} equal to the sum of element means,

$$\tilde{d}_{it} = \tilde{d}_i + \dots + \tilde{d}_t, \quad (3.22)$$

and variance

$$\text{Var}(d_{it}) = \mathbf{1}^T \Sigma \mathbf{1}, \quad (3.23)$$

where $\mathbf{1}$ is an all ones vector in \mathcal{R}^{t-i+1} (Theorem 3.3.2), abbreviated as $d_{it} \sim \mathcal{N}(\tilde{d}_{it}, \mathbf{1}^T \Sigma \mathbf{1})$.

We next introduce the first order loss function $L(x, \omega) = \int_{-\infty}^{\infty} \max(t-x, 0) g_{\omega}(t) d(t)$ and its complementary function $\hat{L}(x, \omega) = \int_{-\infty}^{\infty} \max(x-t, 0) g_{\omega}(t) d(t)$, where ω is a random variable with probability density function $g_{\omega}(\cdot)$, and x is a scalar variable. In what follows, we model the excess back-orders and on-hand stocks in the form of first order loss function and its complementary function.

Let ζ_{it} denote the value of random variable d_{it} , for $t = i, \dots, j$. Since the only replenishment is placed at the beginning of period i , the closing inventory level of period t must 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 excess back-orders of period t in Eq. (3.4) can be reformulated as,

$$\int_{d_i} \dots \int_{d_t} \max(\zeta_{it} - S_i, 0) g_{d_i}(\zeta_i) \dots g_{d_t}(\zeta_t) d(\zeta_i) \dots, d(\zeta_t) = L(S_i, d_{it}), \quad (3.24)$$

and the expected on-hand stocks of period t can be reformulated as,

$$\int_{d_i} \dots \int_{d_t} \max(S_i - \zeta_{it}, 0) g_{d_i}(\zeta_i) \dots g_{d_t}(\zeta_t) d(\zeta_i) \dots, d(\zeta_t) = \hat{L}(S_i, d_{it}). \quad (3.25)$$

Therefore, the expected total cost $C_{ij}(I_{i-1}, S_i)$ over periods i, \dots, j , given initial inventory I_{i-1} , and order-up-to level S_i at the beginning of period i , can be written as follows,

$$C_{ij}(I_{i-1}, S_i) = \sum_{t=i}^j \left(K\delta_t + c \cdot Q_t + h\hat{L}(S_i, d_{it}) + bL(S_i, d_{it}) \right), \quad (3.26)$$

where $\delta_t = \{0, 1\}$, and $Q_t = (S_i - I_{t-1})\delta_t$, for $t = i, \dots, j$. It is clear that, for the single replenishment cycle problem, $\delta_i = 1$, $Q_i = S_i - I_{i-1}$, $\delta_t = 0$, and $Q_t = 0$, for $t = i + 1, \dots, j$.

We now extend the above discussion to a N -replenishment cycle problem over periods i, \dots, j . We assume that the initial inventory level I_{i-1} , the replenishment cycle n , and the corresponding order-up-to levels S_n are fixed at the beginning of period i , where $n = 1, \dots, N$. Therefore, the expected total cost $C_{ij}(I_{i-1}, S_n)$ over periods i, \dots, j is the sum of the expected total cost of each single replenishment cycle n . The order quantity Q_n of replenishment period n is uniquely decided by the opening inventory level I_{n-1} and the order-up-to level S_n , i.e. $Q_n = S_n - I_{n-1}$, for $n = 1, \dots, N$.

Example. We now demonstrate the multi-replenishment cycle problem discussed above on the 4-period problem presented in Section 3.3. We assume that the only two replenishments are placed in period 1 and 3, and the corresponding order-up-to levels are 60 and 100. The resulting expected total cost is 433.88.

3.5.2 An MILP model for computing (R, S) policies

We now present our MILP model for determining optimal (R, S) policies; to approximate the expected holding and penalty cost, we employ the piecewise linear approximation technique proposed by Rossi et al. (2015).

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. (3.24)-(3.25), the expected excess back-order and on-hand stock of period t can be written by means of the first order loss function and its complementary function, $\sum_{j=1}^t L(S_j, d_{jt})P_{jt}$, and $\sum_{j=1}^t \hat{L}(S_j, d_{jt})P_{jt}$. Additionally, since period j must be the most recent order received up to period t , the following constraints must be satisfied.

$$\sum_{j=1}^t P_{jt} = 1, \quad (3.27)$$

$$P_{jt} \geq \delta_j - \sum_{k=j+1}^t \delta_k, \quad j = 1, \dots, t. \quad (3.28)$$

In what follows, let $\tilde{B}_t \geq 0$ and $\tilde{H}_t \geq 0$ denote upper bounds to true values of $\sum_{j=1}^t L(S_j, d_{jt})P_{jt}$, and $\sum_{j=1}^t \hat{L}(S_j, d_{jt})P_{jt}$. We next employ the piecewise linear approximation technique proposed in Rossi et al. (2014, 2015) for d_{jt} to approximate the expected back-orders and on-hand stocks. 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 . Then, the Edmundson-Madansky upper bound can be applied to the expected back-order and on-hand stock.² Therefore, \tilde{B}_t , and \tilde{H}_t are formulated as follows,

$$\tilde{B}_t \geq -\tilde{I}_t + \sum_{j=1}^t S_j P_{jt} \sum_{k=1}^i p_k + \sum_{j=1}^t \left(e_W^{jt} - \sum_{k=1}^i p_k E[d_{jt}|\Omega_i] \right) P_{jt}, \quad (3.29)$$

$$\tilde{H}_t \geq \sum_{j=1}^t S_j P_{jt} \sum_{k=1}^i p_k + \sum_{j=1}^t \left(e_W^{jt} - \sum_{k=1}^i p_k E[d_{jt}|\Omega_i] \right) P_{jt}, \quad (3.30)$$

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

For the special case in which demand follows a standard normal distribution, the piecewise linear approximation parameters p_i , $E[d_{jt}|\Omega_i]$, and e_W^{jt} are provided in Rossi et al. (2014), for $i = 1, \dots, W$, $t = 1, \dots, T$, and $j = 1, \dots, t$. These

²Similarly, the Jensen's lower bound can be applied for approximating the expected excess inventory and back-orders as well, details refer to Rossi. et al. (2015).

parameters can be applied to general normal distributions by using the standardisation formula $\hat{L}(S_j, d_{jt}) = \sigma_{jt} \hat{L}(\frac{S_j - \tilde{d}_{jt}}{\sigma_{jt}}, Z)$ in Rossi et al. (2014), Lemma 7, where σ_{jt} represents the standard deviation of the joint distribution d_{jt} , and Z is a standard normal random variable. Note that, since in our case demand is correlated, the mean \tilde{d}_{jt} and standard deviation σ_{jt} of the demand convolution d_{jt} over periods j, \dots, t must be calculated via Eq. (3.22) and (3.23).

Finally, the expected proportional ordering cost can be reformulated as $c \sum_{t=1}^T Q_t = c\tilde{I}_T + c \sum_{t=1}^T \tilde{d}_t - cI_0$ by adopting the reformulation strategy originally introduced in Tarim and Kingsman (2004) at p. 112. Therefore, the formulation in Fig. 3.1 can be reduced to an equivalent deterministic MILP model given in Fig. 3.2.

$$\min_{\delta_t} -cI_0 + c \sum_{t=1}^T \tilde{d}_t + \sum_{t=1}^T (K\delta_t + h\tilde{H}_t + b\tilde{B}_t) + c\tilde{I}_T \quad (3.31)$$

Subject to, for $t = 1, \dots, T$,

$$\tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1} \geq 0 \quad (3.32)$$

$$\delta_t = 0 \rightarrow \tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1} = 0 \quad (3.33)$$

$$\sum_{j=1}^t P_{jt} = 1, \quad (3.34)$$

$$P_{jt} \geq \delta_j - \sum_{k=j+1}^t \delta_k, \quad j = 1, \dots, t \quad (3.35)$$

$$\tilde{H}_t \geq (\tilde{I}_t + \sum_{j=1}^t \tilde{d}_{jt} P_{jt}) \sum_{k=1}^i p_k + \sum_{j=1}^t \left(e_W^{jt} - \sum_{k=1}^i p_k E[d_{jt} | \Omega_i] \right) P_{jt}, \quad i = 1, \dots, W \quad (3.36)$$

$$\tilde{B}_t \geq -\tilde{I}_t + (\tilde{I}_t + \sum_{j=1}^t \tilde{d}_{jt} P_{jt}) \sum_{k=1}^i p_k + \sum_{j=1}^t \left(e_W^{jt} - \sum_{k=1}^i p_k E[d_{jt} | \Omega_i] \right) P_{jt}, \quad i = 1, \dots, W \quad (3.37)$$

$$\delta_t \in \{0, 1\} \quad (3.38)$$

$$P_{jt} \in \{0, 1\}, \quad j = 1, \dots, t \quad (3.39)$$

Figure 3.2: An MILP model for computing (R, S) policies with correlated demand

The objective (3.31) is to decide the timing and order-up-to level of replenishments so as to minimise the expected total cost comprising ordering, holding, and penalty costs with given initial inventory level I_0 . Constraints (3.32) ensure the non-negativity of replenishments. Constraints (3.33) are indicator constraints (Belotti et al., 2016) capturing the reorder condition. Constraints (3.34) indicate that the most recent replenishment before period t was issued in period j . Constraints (3.35) uniquely define in which period the most recent replenishment prior to t took place. Constraints (3.36)-(3.37) are approximations of expected end of period t holding and penalty costs by means of the first order loss function. Constraints (3.38)-(3.39) set binary variables.

By solving the model in Fig. 3.2, the optimal (R, S) policies are obtained. Specifically, replenishment periods are obtained from δ_t and P_{jt} once and for all,

before demand becomes known. The respective order-up-to levels S_t are obtained by $\tilde{I}_t + \tilde{d}_t$.

Example. We now use the same 4-period example in Section 3.3 to demonstrate the modelling strategy. We solve the MILP model presented in Fig. 3.2, and observe that the review periods are 1 and 3, and corresponding order-up-to levels are 72.15 and 120.01. The simulated expected total cost (ETC) of this policy is 381.75, the computational time required to solve the MILP model is 0.19 seconds. Additionally, in Table 3.1, we compare these results against the optimal policy obtained via SDP.

	ETC	Computational times
MILP	381.75	0.19
SDP	362.55	29.26

Table 3.1: Comparison of the MILP model and the stochastic dynamic program

Our model generalises the discussion in Rossi et al. (2015), which discussed MILP model for approximating optimal (R, S) policy parameters when the demand is independently and identically distributed. Our MILP model exploits the law of total expectation and properties of joint distribution for computing optimal (R, S) policies under correlated demand. As we will show in the next section, our model can be immediately applied to a broad range of time series models drawn from the literature.

3.6 Applications to time-series-based demand processes

In this section we apply the MILP model in Fig. 3.2 for approximating the optimal (R, S) policies with time-series-based demand processes. Our discussion incorporates the AR, MA, ARMA, and AR-ARCH process.

Recall that the MILP model in Fig. 3.2 is built upon properties of the joint distribution of demand convolution d_{jt} over periods j, \dots, t . Once the mean and covariance matrix of demand convolution d_{jt} are decided, the MILP model can be easily implemented and solved by using existing off-the-shelf software such as IBM ILOG Optimisation Studio. Therefore, in this section we mainly focus on presenting the mean, and covariance matrix regarding different time-series-based demand processes.

We first consider the AR process of order P (AR(P)), the MA process of order Q (MA(Q)), and the ARMA process of order P and Q (ARMA(P, Q))

presented in Box et al. (2015). The AR(P) process predicts demand by using past P time periods' demand and a new noise; while the MA(Q) process forecasts demand on previous Q time periods' noises and a new noise. Putting these two processes together yields the complete class of ARMA process of order P and Q (ARMA(P, Q)). For further details please refer to Appendix 3.B. We illustrate associated expressions for means and covariances in Table 3.3.

The AR, MA, and ARCH processes represent the correlations of current demand with realised information. We also consider the AR process with ARCH of order P and M (AR(P)-ARCH(M)) introduced in Engle (1982), where not only the current demand, but also the current noise depends upon the realised information. For further details please refer to Appendix 3.B.4. We present the mean and covariance of an AR(P)-ARCH(M) process in Table 3.3.

Time series	Mean	Covariance
AR(P)	$\frac{\beta_0}{1 - \sum_{p=1}^P \beta_p}$	$\gamma^{ k } = \begin{cases} \sum_{p=1}^P \beta_p \gamma^p + \sigma^2, & k = 0; \\ \sum_{p=1}^P \beta_p \gamma^{ k-p }, & k \geq 1. \end{cases}$
MA(Q)	θ'_0	$\gamma^{ k } = \begin{cases} \sum_{i=0}^{Q-k} \theta_i \theta_{i+k} \sigma^2, & 0 < k \leq Q, \theta_0 = 1; \\ 0, & k > Q. \end{cases}$
ARMA(P, Q)	$\frac{\beta_0}{1 - \sum_{p=1}^P \beta_p}$	N.A.
AR(P)-ARCH(M)	$\frac{\beta_0}{1 - \sum_{p=1}^P \beta_p}$	$\gamma^{ k } = \begin{cases} \sum_{p=1}^P \beta_p \gamma^{ p-k }, & k \geq 1; \\ \sum_{p=1}^P \beta_p \gamma^p + \frac{\alpha_0}{1 - \sum_{m=1}^M \alpha_m}, & k = 0. \end{cases}$

Figure 3.3: Time series processes

3.7 Computational experiments

In this section we present an extensive numerical study to investigate performance of our MILP heuristic discussed in Section 3.5.2. We first design a test bed featuring instances defined over an 8-period planning horizon in Section 3.7.1. On this test bed, we assess the behaviour of the optimality gap and the computational efficiency of our MILP heuristic on multivariate normally distributed demand. We then assess the computational performance of our MILP model on time-series-based demand processes over a 15-period planning horizon in Section 3.7.2. Numerical experiments are conducted by using IBM ILOG CPLEX Optimization Studio 12.7 and MATLAB R2016a on a 3.2GHz Intel(R) Core(TM) with 8GB of RAM.

3.7.1 Multivariate normal distribution

We consider a test bed which includes 320 instances. Specifically, we consider ten general multivariate normal distributed demand patterns displayed in Fig. 3.4,

comprising two life cycle patterns (LCY1 and LCY2), two sinusoidal patterns (SIN1 and SIN2), a stationary pattern (STA), a random pattern (RAND), and four empirical patterns (EMP1, ..., EMP4). Full details on the experimental set-up are given in Appendix 3.C. We assume that the current demand is only related to the past one period demand with covariance coefficient $\rho = \{0.25, 0.5\}$. The fixed ordering cost K ranges in $\{200, 400\}$, the proportional ordering cost c ranges in $\{0, 1\}$, and the penalty cost b takes values $\{10, 20\}$. The proportional holding cost $h = 1$. We further assume that the demand is normally distributed with coefficient of standard deviation $c_v = \{0.15, 0.3\}$ (note that $\sigma_{d_t} = c_v \cdot \mu_t$). Since we operate under the assumption of normality, our models can be readily linearised by using the piecewise linearisation parameters available in Rossi et al. (2014).

We utilize the stochastic dynamic program (SDP) model discussed in Section 3.3 as a benchmark. We compare against this benchmark in terms of optimality gap and computational time. We first obtain optimal parameters for each test instance by implementing the SDP algorithm in Matlab. We then solve each test instance by implementing the MILP model in IBM ILOG CPLEX Optimization Studio. Specifically, for the MILP model, we employ eleven segments in the piecewise-linear approximations of B_t and H_t (for $t = 1, \dots, T$) in order to guarantee reasonable computational performance. To estimate the cost of the policies obtained via our heuristics, we simulate all policies via Monte Carlo Simulation (100,000 replications).

Table 3.2 gives an overview of optimality gaps (%) of the MILP model discussed in Section 3.5.2 for different pivoting parameters. The optimality gap is defined as the difference between the simulated expected total cost obtained via the MILP model and the SDP model. We observe that it is difficult to make a general remark on the demand patterns. An increase of fixed ordering cost, proportional ordering cost, and covariance coefficient slightly reduces the optimality gap; conversely, an increase of penalty cost and coefficient of variation tends to increase the optimality gap. More specifically, when the fixed ordering cost increases from 200 to 300, the optimality gap decreases from 2.59% to 1.99%; while the optimality gap increases from 0.98% to 3.58% as the coefficient of variation increases from 0.15 to 0.3. On average, the optimality gap of the MILP heuristic on the multivariate demand is 2.28%.

We assess the accuracy of the MILP model by comparing the cost predicted by our model against the cost obtained via simulation in Table 3.2. We notice that the average model accuracy is 0.47%.




Figure 3.4: Demand patterns in our computational analysis

	MILP		RHC
	MILP gap(%)	MILP accuracy (%)	RHC gap(%)
Demand pattern			
LCY1	4.88	0.72	3.16
LCY2	1.30	0.33	0.90
SIN1	1.44	0.49	0.99
SIN2	1.51	0.50	1.12
RAND	1.09	0.31	0.95
STA	1.97	0.51	1.49
EMP1	3.81	0.49	1.94
EMP2	2.71	0.42	0.74
EMP3	0.63	0.34	0
EMP4	3.45	0.55	1.67
Fixed ordering cost			
200	2.57	0.50	1.31
300	1.99	0.43	1.27
Proportional ordering cost			
0	2.38	0.50	1.37
1	2.18	0.43	1.21
Penalty cost			
10	1.86	0.33	1.11
20	2.70	0.60	1.47
Coefficient of variation			
0.15	0.98	0.35	0.72
0.3	3.58	0.58	1.86
Covariance coefficient			
0.25	2.30	0.46	1.61
0.5	2.26	0.48	0.98
Average gap	2.28	0.47	1.29

Table 3.2: Computational behaviours of the multivariate normally distributed demand processes for different pivoting parameters

We also investigate a receding horizon control (RHC) implementation (Kwon and Han, 2006), in which we repeatedly solve the lot-sizing problem at each time period to determine an up-to-date optimal plan which takes into account all information available. We use Monte Carlo simulation to estimate the expected total cost; as a stopping criterion we impose a maximum estimation error of 0.03% of the estimated cost at 95% confidence. In Table 3.2 we present the RHC gap³ and the differences between the expected total cost obtained via RHC and SDP. We observe that the average RHC gap is 1.29%.

We present computational times of both the MILP model discussed in Section 3.5.2 and SDP model in Table 3.3. Note that “STDEV” in Table 3.3 represents the standard deviation. We observe that the computational time of the MILP model is not significantly affected by the demand patterns; while that of the SDP algorithm fluctuates widely. Furthermore, we observe that the fixed ordering cost, proportional ordering cost, penalty cost, coefficient of variation, and covariance coefficient do not have significant effect on the computational efficiency of both the MILP and SDP algorithm. In general, the average computational time of the SDP and MILP algorithm are 192.17 and 0.10 seconds; their standard deviations are 164.89 and 0.10 seconds.

3.7.2 Time-series-based demand processes

In this section we demonstrate that the MILP algorithm discussed in Section 3.5.2 can also be extended to solve lot-sizing problems with time-series-based demand. Existing algorithms in the literature can only tackle lower order AR, MA, or ARMA processes; in what follows we will show that higher order time-series-based demand processes are tractable with our MILP model.

We only assess the model accuracy and computational efficiency of the MILP algorithm on the time-series-based demand processes, since time series processes are built upon the multivariate normal distribution and we have already investigated optimality gaps in Section 3.7.1 (on average, 2.28%). Additionally, using SDP to tackle higher order time-series-based demand processes is computationally prohibitive even for very small instances.

We consider a 15-period test bed which includes 112 instances. Specifically, we assess the computational performance of the MILP model on eight different time-series-based processes, comprising AR(1), AR(3), MA(1), MA(3), ARMA(1, 1), ARMA(3, 3), AR(1)-ARCH(1), and AR(3)-ARCH(3). These time-series-based

³Computational experiments are conducted by using the IBM ILOG CPLEX Optimization Studio 12.7 and Eclipse 4.7.3 on a 1.2GHz Intel Core i5 with 4GB 1600MHz DDR3.

Settings	SDP		MILP	
	Mean	STDEV	Mean	STDEV
Demand pattern				
LCY1	65.37	1.92	0.10	0.07
LCY2	55.76	1.10	0.08	0.04
SIN1	55.49	0.96	0.12	0.17
SIN2	55.02	0.35	0.08	0.04
RAND	57.69	0.25	0.12	0.07
STA	56.11	1.18	0.10	0.06
EMP1	394.66	7.20	0.17	0.22
EMP2	393.21	2.34	0.08	0.03
EMP3	394.16	1.26	0.09	0.03
EMP4	394.26	1.32	0.10	0.03
Fixed ordering cost				
200	192.12	164.83	0.10	0.06
300	192.22	164.96	0.11	0.13
Proportional ordering cost				
0	192.22	165.15	0.11	0.07
1	192.13	164.81	0.10	0.13
Penalty cost				
10	192.53	165.39	0.12	0.13
20	191.81	164.39	0.09	0.05
coefficient of variation				
0.15	192.02	164.48	0.11	0.09
0.3	181.73	165.30	0.10	0.11
covariance coefficient				
0.25	191.96	164.55	0.11	0.12
0.5	192.39	165.24	0.10	0.07
Average	192.17	164.89	0.10	0.10

Table 3.3: Computational times of the multivariate normally distributed demand processes for different pivoting parameters

processes are generated with expressions in Fig. 3.5. We assume that the coefficient of variance $cv = \{0.15, 0.3\}$, the fixed ordering cost $K = \{200, 300\}$, the proportional ordering cost $c = \{0, 1\}$, the holding cost $h = 1$, and the penalty cost $b = \{10, 20\}$.

- AR(1): $d_t = 25 + 0.75d_{t-1} + \epsilon_t$, $\epsilon_t \sim \mathcal{WN}(0, (cv * \tilde{d}_t)^2)$;
- AR(3): $d_t = 25 + 0.5d_{t-1} + 0.2d_{t-2} + 0.1d_{t-3} + \epsilon_t$, $\epsilon_t \sim \mathcal{WN}(0, (cv * \tilde{d}_t)^2)$;
- MA(1): $d_t = 100 + 0.75\epsilon_{t-1} + \epsilon_t$, $\epsilon_t \sim \mathcal{WN}(0, (cv * \tilde{d}_t)^2)$;
- MA(3): $d_t = 100 + 0.5\epsilon_{t-1} + 0.2\epsilon_{t-2} + 0.1\epsilon_{t-3} + \epsilon_t$, $\epsilon_t \sim \mathcal{WN}(0, (cv * \tilde{d}_t)^2)$;
- ARMA(1, 1): $d_t = 25 + 0.75d_{t-1} + 0.75\epsilon_{t-1} + \epsilon_t$, $\epsilon_t \sim \mathcal{WN}(0, (cv * \tilde{d}_t)^2)$;
- ARMA(3, 3): $d_t = 25 + 0.5d_{t-1} + 0.2d_{t-2} + 0.1d_{t-3} + 0.5\epsilon_{t-1} + 0.2\epsilon_{t-2} + 0.1\epsilon_{t-3} + \epsilon_t$, $\epsilon_t \sim \mathcal{WN}(0, (cv * \tilde{d}_t)^2)$;
- AR(1)-ARCH(1): $d_t = 25 + 0.75d_{t-1} + \epsilon_t$, $\epsilon_t = \mu_t \sqrt{100 + 0.75\epsilon_{t-1}^2}$, where $\mu_t \sim \mathcal{IIN}(0, 1)$;
- AR(3)-ARCH(3): $d_t = 25 + 0.5d_{t-1} + 0.2d_{t-2} + 0.1d_{t-3} + \epsilon_t$, $\epsilon_t = \mu_t \sqrt{1 + 0.3\epsilon_{t-1}^2 + 0.2\epsilon_{t-2}^2 + 0.1\epsilon_{t-3}^2}$, where $\mu_t \sim \mathcal{IIN}(0, 1)$.

Figure 3.5: Expressions for time-series-based demand patterns

Table 3.4 demonstrates computational times of the MILP algorithm discussed in Section 3.5.2 for different pivoting parameters.⁴ It is difficult to draw a general remark on different demand patterns. We observe that an increase of fixed ordering cost slightly increases the computational time; while the increase of proportional ordering cost, penalty cost, and coefficient of variation decreases the computational time. For instance, the computational time increases from 0.59 to 0.77 seconds as the fixed ordering cost increases from 200 to 300; Additionally, when the proportional ordering cost increases from 0 to 1, the average computational time drops from 0.77 to 0.59 seconds. On average the computational time is 0.68 seconds, and the standard deviation is 0.26 seconds.

Table 3.5 presents the model accuracy of the MILP algorithm on time-series-based demand processes. We observe that the average model accuracy is 3.41%.

⁴Please note that the coefficient of variation has no effect on the AR-ARCH processes, therefore, only AR, MA, and ARMA processes are considered in doing sensitivity analysis.

Settings	Computational time (s)	
	Average	STDEV
Demand patterns		
AR(1)	0.63	0.23
AR(3)	0.63	0.15
MA(1)	0.90	0.20
MA(3)	0.69	0.10
ARMA(1,1)	0.69	0.41
ARMA(3,3)	0.51	0.35
AR(1)-ARCH(1)	0.75	0.28
AR(3)-ARCH(3)	0.68	0.34
Fixed ordering costs		
200	0.59	0.30
300	0.77	0.27
Proportional ordering costs		
0	0.77	0.35
1	0.59	0.20
Penalty ordering costs		
10	0.72	0.26
20	0.64	0.33
Coefficient of variation		
0.15	0.72	0.24
0.3	0.63	0.36
Average gap	0.68	0.26

Table 3.4: Computational efficiency of the time-series-based demand processes for different pivoting parameters

Settings	Model accuracy (%)
Demand patterns	
AR(1)	7.41
AR(3)	2.50
MA(1)	4.10
MA(3)	0.67
ARMA(1,1)	2.50
ARMA(3,3)	3.53
AR(1)-ARCH(1)	6.30
AR(3)-ARCH(3)	0.24
Fixed ordering costs	
200	3.39
300	3.46
Proportional ordering costs	
0	3.91
1	2.95
Penalty ordering costs	
10	2.86
20	4.00
coefficient of variation	
0.15	3.03
0.3	3.96
Average gap	3.41

Table 3.5: Model accuracy of the time-series-based demand processes for different pivoting parameters

3.8 Conclusion

In this paper, we consider a single-item single-stock location inventory lot-sizing problem under non-stationary stochastic correlated demand, fixed and unit ordering cost, holding cost, and penalty cost. We present an MILP-based model for approximating optimal (R, S) policies under normally distributed demand featuring correlation across periods as well as under a collection of time-series-based demand processes. In contrast to other approaches in the literature, our model can be easily implemented and solved by using off-the-shelf mathematical programming packages such as IBM ILOG optimisation studio.

We conducted an extensive numerical study comprising 432 instances. We first investigated the behaviours of the optimality gap and computational efficiency of the MILP heuristic on a 8-period test bed with 320 instances. We observe that the optimality gap is 2.28%, and the average computational time is 0.1 seconds.

We then assessed the computational efficiency of the proposed MILP model on time-series-based demand processes over a 15-period planning horizon comprising 112 instances. We observe that the average computational time is 0.68 seconds.

Our computational study demonstrates that our model is computationally efficient and accurate. Moreover, in contrast to existing approaches in the literature, it can tackle higher order time-series-based demand processes.

Appendix

3.A List of symbols

In this section we present a list of symbols used in this paper.

T	periods in the planning horizon
d_t	random variable
$g(\cdot)$	probability density function
d	a vector
η_t	realisation of random variable d_t
\mathcal{D}_t	realised demand set at the beginning of period t
ζ_t	value of random variable d_t
I_t	inventory level at the end of period t
Q_t	ordering quantity placed at the beginning of period t
I_0	initial inventory level at the beginning of the planning horizon
$c(\cdot)$	ordering cost

K	fixed ordering cost
c	proportional ordering cost
h	proportional holding cost
b	proportional penalty cost
\mathcal{S}_t	system state at the beginning of period t , $\mathcal{S}_t = \{I_{t-1}, \mathcal{D}_t\}$
$f_t(I_{t-1}, \mathcal{D}_t, Q_t)$	immediate cost of period t with opening inventory level I_{t-1} , realised demand set \mathcal{D}_t , and order quantity Q_t
$C_t(I_{t-1}, \mathcal{D}_t)$	the expected total cost of an optimal policy over period t, \dots, T with opening inventory level I_{t-1} and realised demand set \mathcal{D}_t
S_t	order-up-to level of period t
δ_t	binary variable
$\hat{C}_t(I_{t-1}, \mathcal{D}_t)$	expected total cost over periods t, \dots, T under the “static-dynamic uncertainty” policy with opening inventory level I_{t-1} and realised demand set \mathcal{D}_t
M	a large number
$\check{C}_t(I_{t-1}, \mathcal{D}_t)$	expected total cost over period t, \dots, T under (R, S) policy with opening inventory level I_{t-1} and realised demand set \mathcal{D}_t
\tilde{d}_t	the expected value of random variable d_t
d_{jt}	a random variable denotes the demand over period j, \dots, t , i.e. $d_{jt} = d_j + \dots + d_t$
\tilde{d}_{jt}	expected value of the convolution $\tilde{d}_j + \dots + \tilde{d}_t$
ζ_{jt}	value of random variable d_{jt}
ω	a random variable
x	a scalar value
$L(x, \omega)$	first order loss function
$\hat{L}(x, \omega)$	complementary of first order loss function
$C_{ij}(I_{i-1}, S_i)$	the expected total cost over periods i, \dots, j with opening inventory level I_{i-1} and order-up-to level S_i at the beginning of period i
P_{jt}	a binary variable 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
Ω	support of d_{jt}
W	number of regions in a partition of Ω
i	region index ranging in $1, \dots, W$

Ω_i	the i^{th} subregion of Ω
p_i	$Pr(d_{jt} \in \Omega_i)$
$E[d_{jt} \Omega_i]$	conditional expectation of d_{jt} in Ω_i
\tilde{H}_t	the upper bound to the true value of $\sum_{j=1}^t \hat{L}(S_j, d_{jt})P_{jt}$
\tilde{B}_t	the upper bound to the true value of $\sum_{j=1}^t L(S_j, d_{jt})P_{jt}$
e_W^{jt}	approximation error
σ_{jt}	the standard deviation of d_{jt}
Z	a standard normal random variable

Table 3.6: A list of symbols

3.B Time series processes

In this section we present the AR, MA, ARMA, and AR-ARCH processes.

3.B.1 AR process

An AR process operates under the assumption that there is some linear correlation between values in a time series.

Definition 3.B.1. (Autoregressive process of Order P). Consider a random variable d_t , $t = \{1, \dots, T\}$, the AR process of order P , abbreviated AR(P), is defined by the equation

$$d_t = \beta_0 + \sum_{p=1}^P \beta_p d_{t-p} + \epsilon_t, \quad \text{where } \{\epsilon_t\} \sim \mathcal{N}(0, \sigma^2) \quad (3.40)$$

where $\beta_0, \beta_1, \dots, \beta_P$ are parameters of this model, and $\{\epsilon_t\}$ is a sequence of normally distributed independent random variables with mean 0 and variance σ^2 .

AR(P) process has the following properties.

- Since the AR(P) is a weakly stationary process, it has constant mean $E[d_t]$, $E[d_t] = \frac{\beta_0}{1 - \sum_{p=1}^P \beta_p}$;
- Let $\text{Var}(d_t)$ denote the variance of the AR(P) process, $\text{Var}(d_t) = \frac{\sigma^2}{1 - \sum_{p=1}^P \beta_p^2}$;
- Let $\gamma^{|k|}$ be the covariance of d_t with itself at a different point in time, as the k^{th} auto-covariance. Then,

$$\gamma^{|k|} = \begin{cases} \sum_{p=1}^P \beta_p \gamma_p + \sigma^2, & k = 0; \\ \sum_{p=1}^P \beta_p \gamma^{|k-p|}, & \text{for } |k| \geq 1. \end{cases} \quad (3.41)$$

With the properties presented above and Theorem 3.3.2, the mean and covariance matrix of the demand convolution d_{jt} over period j, \dots, t are pre-computed. Therefore, the stochastic lot-sizing problems with AR(P) demand process can be easily adjusted and solved with the MILP model in Fig. 3.2. This problem is also resolvable via stochastic dynamic programming. However, since the current demand is linearly correlated to past P periods, the stochastic dynamic programming formulation is complex and hard to solve.

We note that, this AR demand process permits negative demand. However, in most industrial contexts, negative demand is unlikely or not allowed. Hence, as with any model, some judgement is required as to the applicability of this model of the demand process to the real world.

3.B.2 MA process

Regarding the AR(P) demand processes, the demand in period t depends on the realised demand of last P periods, while its realised noises have no effects on its current value. In this session we present the MA process where the current demand depends on only current noise and realised noises instead of demand.

Definition 3.B.2. (Moving Average process of Order Q). An MA process of order Q (MA(Q)) has dynamics which follows

$$d_t = \theta'_0 + \sum_{q=1}^Q \theta_q \epsilon_{t-q} + \epsilon_t, \quad \text{where } \{\epsilon_t\} \sim \mathcal{N}(0, \sigma^2) \quad (3.42)$$

where $\theta'_0, \theta_1, \dots, \theta_Q$ are parameters of this model, and $\{\epsilon_t\}$ is a sequence of normally distributed independent random variables with mean 0 and variance σ^2 .

The MA(Q) process has the following properties.

- Since the MA(Q) is a weakly stationary process, it has constant process mean $E[d_t], E[d_t] = \theta'_0$;
- Let $\text{Var}(d_t)$ denote the variance of the MA(Q) process, $\text{Var}(d_t) = (1 + \sum_{q=1}^Q \theta_q^2) \sigma^2$;
- Let $\gamma^{|k|}$ be the covariance of d_t with itself at a different point in time, as the k^{th} auto-covariance. Then,

$$\gamma^{|k|} = \begin{cases} \sum_{i=0}^{Q-k} \theta_i \theta_{i+k} \sigma^2, & 0 < |k| \leq Q, \text{ where } \theta_0 = 1; \\ 0, & |k| > Q. \end{cases} \quad (3.43)$$

3.B.3 ARMA process

The AR process stipulates that the current value depends on its previous values and a new noise; while the current value in a MA process depends on both a new noise and previous noises. By putting these two processes together yields the complete class of ARMA process.

Definition 3.B.3. (Autoregressive Moving Average process of Order P and Q). An ARMA process with orders P and Q (ARMA(P,Q)) is defined as follows,

$$d_t = \beta_0 + \sum_{p=1}^P \beta_p d_{t-p} + \sum_{q=1}^Q \theta_q \epsilon_{t-q} + \epsilon_t, \quad \text{where } \{\epsilon_t\} \sim \mathcal{N}(0, \sigma^2) \quad (3.44)$$

Where $\beta_0, \beta_1, \dots, \beta_P$, and $\theta_1, \dots, \theta_Q$ are parameters of this model, and $\{\epsilon_t\}$ is a sequence of normally distributed independent random variables with mean 0 and variance σ^2 .

The ARMA(P, Q) process has constant mean $E[d_t] = \frac{\beta_0}{1 - \sum_{p=1}^P \beta_p}$, while the variance $\text{Var}(d_t)$ cannot be easily expressed. We take the ARMA(1,1) process as an example to show its mean, variance, and auto-covariance. The ARMA(1,1) is defined as,

$$d_t = \beta_0 + \beta_1 d_{t-1} + \theta_1 \epsilon_{t-1} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \sigma^2). \quad (3.45)$$

It has the following properties.

- Since the ARMA(1,1) is a weakly stationary process, it has constant process mean $E[d_t] = \frac{\beta_0}{1 - \beta_1}$;
- Let $\text{Var}(d_t)$ denote the variance of ARMA(1,1) process, $\text{Var}(d_t) = \frac{1 + \theta_1^2 + 2\theta_1\beta_1}{1 - \beta_1^2} \sigma^2$;
- Let $\gamma^{|k|}$ be the covariance of d_t with itself at a different point in time, as the k^{th} auto-covariance. Then,

$$\gamma^{|k|} = \begin{cases} \beta_1 \gamma^{|k|-1} + \theta_1 \sigma^2, & |k| = 1 \\ \beta_1 \gamma^{|k|-1}, & |k| \geq 2 \end{cases} \quad (3.46)$$

3.B.4 AR-ARCH process

We have discussed in previous sections that the AR, MA, and ARMA processes represent the correlations of current demand with realised information. This section will present a class of models where not only the current demand, but

also the current noise depend upon the realised information. We first present the linear ARCH process originally introduced by Engle (1982), the time varying conditional variance is postulated to be a linear function of the past M squared innovations. We further present the AR-ARCH process.

Definition 3.B.4. (Autoregressive Conditional Heteroskedasticity process) An ARCH process of order M (ARCH(M)) has dynamics which follows

$$\epsilon_t = \mu_t \sqrt{\alpha_0 + \sum_{m=1}^M \alpha_m \epsilon_{t-m}^2}, \quad \mu_t \sim \mathcal{N}(0, 1) \quad (3.47)$$

where $\alpha_0, \alpha_1, \dots, \alpha_M$ are positive parameters.

Let F_t denote the information set available at time t . The conditional mean of ϵ_t is $E[\epsilon_t|F_t] = E[\mu_t|F_t] \cdot \sqrt{\alpha_0 + \sum_{m=1}^M \alpha_m \epsilon_{t-m}^2} = 0$ since $E[\mu_t|F_t] = 0$. The conditional variance of ϵ_t is $\text{Var}(\epsilon_t|F_{t-1}) = E[\text{Var}(\epsilon_t^2|F_{t-1})] - E[\epsilon_t|F_{t-1}]^2 = \alpha_0 + \sum_{m=1}^M \alpha_m \epsilon_{t-m}^2$.

Therefore, the unconditional mean of ϵ_t is $E[\epsilon_t] = 0$. The unconditional variance of ϵ_t is $\text{Var}(\epsilon_t) = E[\epsilon_t^2|F_{t-1}] + \text{Var}(E[\epsilon_t|F_{t-1}]) = \alpha_0 + \sum_{m=1}^M \alpha_m \text{Var}(\epsilon_{t-m})$. Since the ARCH(M) is a stationary process, $\text{Var}(\epsilon_t) = \frac{\alpha_0}{1 - \sum_{m=1}^M \alpha_m}$.

An AR(P) process has an ARCH-free white-noise process $\{\epsilon_t\}$ with variance σ^2 . If we assume that the white noise process $\{\epsilon_t\}$ now is a ARCH(M) process, we have a more complicated AR(P) process with ARCH(M) effects.

Definition 3.B.5. (Autoregressive process with Autoregressive Conditional Heteroskedasticity effects). An AR process of order P with ARCH of order M (AR(P)-ARCH(M)) has dynamics which follows

$$d_t = \beta_0 + \sum_{p=1}^P \beta_p d_{t-p} + \epsilon_t, \quad \epsilon_t = \mu_t \cdot \sqrt{\alpha_0 + \sum_{m=1}^M \alpha_m \epsilon_{t-m}^2} \quad (3.48)$$

where $\mu_t \sim \mathcal{N}(0, 1)$, $\beta_0, \beta_1, \dots, \beta_P$, and $\alpha_0, \alpha_1, \dots, \alpha_M$ are parameters of this model. Additionally, to ensure that $\epsilon_{\epsilon_t} \geq 0$, we need $\alpha_0 > 0$, and $\alpha_m \geq 0$ for $m \in \{1, \dots, M\}$.

Like the AR(P) process, the AR(P)-ARCH(M) process is a weakly stationary process, it has stationary mean, variance, and covariance as follows.

- The unconditional mean of the AR(P)-ARCH(M) is, $E[d_t] = \frac{\beta_0}{1 - \sum_{p=1}^P \beta_p}$;
- the variance is, $\text{Var}(d_t) = \frac{\alpha_0}{(1 - \sum_{p=1}^P \beta_p^2)(1 - \sum_{m=1}^M \alpha_m)}$;

- The k^{th} auto-covariance is,

$$\gamma^{|k|} = \begin{cases} \sum_{p=1}^P \beta_p \gamma^{|p-k|}, & |k| \geq 1; \\ \sum_{p=1}^P \beta_p \gamma^p + \frac{\alpha_0}{1 - \sum_{m=1}^M \alpha_m}, & k = 0. \end{cases} \quad (3.49)$$

3.C Test bed

Periodic demand with different demand patterns under our eight period computational study are displayed in Table 3.7. The first column represents period indices; other columns represent demand patterns.

Period	LCY1	LCY2	SIN1	SIN2	STA	RAND	EMP1	EMP2	EMP3	EMP4
1	15	3	15	12	10	2	5	4	11	18
2	16	6	4	7	10	4	15	23	14	6
3	15	7	4	7	10	7	26	28	7	22
4	14	11	10	10	10	3	44	50	11	22
5	11	14	18	13	10	10	24	39	16	51
6	7	15	4	7	10	10	15	26	31	54
7	6	16	4	7	10	3	22	19	11	22
8	3	15	10	12	10	3	10	32	48	21

Table 3.7: Expected demand data of the 8-period computational analysis

Chapter 4

Paper III: Nonstationary (R, S) policies for joint replenishment inventory systems

Abstract

This paper considers the periodic-review nonstationary stochastic joint replenishment problem (JRP) under Bookbinder and Tan's static-dynamic uncertainty control policy. According to a static-dynamic uncertainty control rule, the decision maker fixes timing of replenishments once and for all at the beginning of the planning horizon, inventory position is then raised to a predefined order-up-to position at the beginning of each replenish period. We present a mixed integer linear programming (MILP) model for approximating optimal static-dynamic uncertainty policy parameters. We further demonstrate that our MILP model can be used to approximate the optimal control rule for the JRP, also known as (σ, \vec{S}) policy. An extensive computational study illustrates the effectiveness of our approach when compared to other competitor approaches in the literature.

4.1 Introduction

The Joint Replenishment Problem (JRP) occurs when several items are ordered from the same supplier, or several products have the same means of transportation, or several products are processed on the same piece of equipment (Salameh

et al., 2014). Every time an order is placed, the group fixed ordering cost is incurred regardless the number of items replenished; in addition there are also item-specific fixed ordering costs that are charged whenever an item is included in a replenishment order. The goal of the JRP is to determine the optimal inventory replenishment plan that minimises the cost of replenishing multiple items.

Literature on JRP can be roughly categorised into deterministic and stochastic based on the nature of demand. In the deterministic joint replenishment inventory system, demand for each individual item is assumed to be known over an infinite time horizon and replenishments are made at equally spaced time intervals; the problem is to determine the length of replenishment cycles and the frequency of replenishing individual items, e.g., (Goyal and Belton, 1979; Kaspi and Rosenblatt, 1991; Viswanathan, 1996; Wildeman et al., 1997; Hariga, 1994; Goyal and Deshmukh, 1993; Boctor et al., 2004; Nilsson et al., 2007). In the stochastic joint replenishment inventory system, the demand for each individual item is unknown, but follows a certain type of distribution; the problem is to decide the optimal parameters of a given inventory policy, e.g., (Balintfy, 1964; Atkins and Iyogun, 1988; Renberg and Planche, 1967; Kalpakam and Arivarignan, 1993; Viswanathan, 1997; Nielsen and Larsen, 2005; Özkaya et al., 2006). Most literature still presents applications to deterministic demand; however, the study regarding stochastic demand has received increasing attention due to its practical relevance (Bastos et al., 2017). This work belongs to the growing literature on the stochastic joint replenishment.

This paper applies the static-dynamic strategy, proposed by Bookbinder and Tan (1988) for tackling single-item lot-sizing problems, in the context of a JRP system. The static-dynamic strategy, known as (R, S) , features two control parameters: R , timing of replenishment, and S , order-up-to position. At each review period, the decision maker places an order so as to increase the inventory position (net inventory level + outstanding orders) to a given order-up-to position. In the context of the JRP system, a periodic-review (R, S) policy is adopted for each item. The (R, S) policy is an appealing strategy since it eases the coordination between supply chain players (Kilic and Tarim, 2011), and facilitates managing joint replenishment (Silver et al., 1998).

Our goal is to tackle the periodic-review stochastic JRP under (R, S) policy. We first present a mixed-integer linear programming (MILP) model for computing optimal policy parameters that optimise the expected total cost comprising group fixed ordering costs, item-specific ordering costs, holding costs, and penalty costs over the planning horizon. Our model generalises Rossi et al. (2015), which

discussed an MILP model for approximating (R, S) policy parameters for single-item lot-sizing problems. We further show that our MILP model can be used to approximate (σ, \vec{S}) policies, which are known to be optimal for this class of problem (Liu and Esogbue, 2012). Under this policy, decision makers order up to \vec{S} if opening inventory positions fall in σ ($\vec{S} \in \mathcal{R}^N$, N represents the number of items) at the beginning of each time period. The definition of σ is general; its shape and properties are literally unknown. There is no guarantee of σ by convex, or even connected. Numerical experiments illustrate the effectiveness of our models.

We contribute to the literature on the stochastic JRP as follows.

- We present an MILP model for tackling the nonstationary stochastic JRP under (R, S) policy.
- We demonstrate that the MILP model can be used to approximate (σ, \vec{S}) policies.
- In an extensive computational study based on existing test beds drawn from the literature we demonstrate the effectiveness of our models when compared to other competing approaches in the literature.

The rest of this paper is organised as follows. Section 4.2 surveys relevant literature. Section 4.3 describes problem settings. Section 4.4 presents an MILP model for computing (R, S) policy parameters. Section 4.5 extends the MILP model for approximating the optimal (σ, \vec{S}) policy parameters. An extensive computational study is conducted in Section 4.6. We draw conclusions in Section 4.7.

4.2 Literature review

The problem of controlling the inventory of a multi-item system under joint replenishment has received increasing attention over the past several decades. For a thorough review of literature readers could refer to (Silver and Peterson, 1985; Goyal and Satir, 1989; Van Eijs et al., 1992; Khouja and Goyal, 2008; Bastos et al., 2017). In this section, we focus our attention on existing policies for tackling stochastic JRPs. In particular, we survey control policies that have been considered in the literature.

(σ, \vec{S}) policy. Since the landmark study Scarf (1960) proved the optimality for the single-item inventory problem, there have been several attempts to prove the

optimality for multi-item inventory systems. Johnson (1967) proved the optimal policy in the stationary case is a (σ, \vec{S}) policy, where $\sigma \subset \mathcal{R}^N$ and $\vec{S} \in \mathcal{R}^N$, and one orders up to \vec{S} if inventory levels $\vec{I} \in \sigma$ and $\vec{I} \leq \vec{S}$ and one does not order if $\vec{I} \notin \sigma$. Kalin (1980) showed when $\vec{I} \in \sigma$ and $\vec{I} \not\leq \vec{S}$, there exists $\vec{S}(\vec{I}) \geq \vec{I}$ such that the optimal policy is to order up to $\vec{S}(\vec{I})$, this policy is named $(\sigma, \vec{S}(\cdot))$ policy. Ohno and Ishigaki (2001) proved the optimality of $(\sigma, \vec{S}(\cdot))$ policy for continuous-time inventory problems with compound Poisson demand. Gallego and Sethi (2005) gave the general definition of K -convexity in \mathcal{R}^N , which encompasses both the joint ordering and individual ordering case.

(s, c, S) policy. Several works on stochastic JRPs have focused on computing (s, c, S) policies, introduced by Balintfy (1964). This policy features three control parameters: s , reorder point; c , can-order level; S , order-up-to position. Under this policy, decision makers order up to S when either at a demand epoch the inventory position drops to or below s ; or when at a special replenishment opportunity the inventory position is at or below c . Under the assumption of Poisson-distributed demand, Ignall (1969) proved that the (s, c, S) policy is not optimal even for two-item problems. Silver (1974) proposed the decomposition method to compute (s, c, S) policy parameters, where the multi-item problem is decomposed into several single-item problems. This approximation technique was followed by (Melchiors, 2002; Johansen and Melchiors, 2003). Kayış et al. (2008) modelled the two-item JRP problem as a semi-Markov decision model, and proposed an enumerative approach to approximate (s, c, S) policies. In addition, (Schaack and Silver, 1972; Thompstone and Silver, 1975; Silver, 1981; Federgruen et al., 1984) studied JRPs with compound Poisson-distributed demand.

(R, T) policy. Atkins and Iyogun (1988) proposed two periodic-review (R, T) -type policies, namely periodic policy P and modified periodic policy MP , which differ only in the way the ordering periods T_i are determined. Under this policy, every T_i periods, the inventory position of item i is raised to R_i . Numerical experiments demonstrate that the MP policy performs consistently better than the (s, c, S) policy, and the P policy generally outperforms the (s, c, S) policy excepting problems involving small values of group fixed ordering cost.

(Q, S) policy. This policy was first proposed by Renberg and Planche (1967). Under this policy, whenever the total inventory position drops to the group reorder point, an order is placed to raise inventory position of each item to item-specific order-up-to position S . The combined order quantity is Q , and the group reorder point is reached when the combined usage reaches Q . Pantumsinchai (1992) evaluated the computational performance of the (Q, S) policy by compar-

ing it against the (s, c, S) policy, P policy and MP policy on the basis of long-run total average costs. Computational experiments showed that the MP policy consistently outperforms the (s, c, S) policy on the test instances, and both MP and (Q, S) policy perform better as the group ordering cost increases. The study showed that the (Q, S) policy is appropriate for items for which the stock-out costs are low and the major set-up cost is high relative to the minor set-up cost.

$P(s, S)$ policy. This policy was proposed by Viswanathan (1997) for periodic-review inventory systems, in which inventory position of each item is reviewed at every fixed and constant time interval. At each review time, the (s, S) policy is applied to each item, so that any item with inventory position at or below s is ordered up to S . For a fixed review period, the algorithm of Zheng and Federgruen (1991) is adopted to compute the optimal (s, S) policy parameters. Computational studies indicated that although the proposed policy requires more computational effort, it generally dominates the MP policy, and dominates (s, c, S) policy, and (Q, S) policy for most test instances.

$Q(s, S)$ policy. Nielsen and Larsen (2005) combined features of (Q, S) policy and $P(s, S)$ policy, and proposed the $Q(s, S)$ policy. By operating under this policy, the total inventory position is continuously reviewed while the item-specific inventory positions are reviewed only when the total consumption since the last order reaches Q . Then every item with inventory position less than or equal to its respective reorder point s is ordered to S . An analytic solution is derived by using the Markov decision theory in Nielsen and Larsen (2005). Computational study demonstrated that the $Q(s, S)$ policy outperforms $P(s, S)$ policy, and dominates (Q, S) policy in 17 of 18 test instances on the data set of Atkins and Iyogun (1988).

(Q, S, T) policy. This continuous-review policy was proposed by Özkaya et al. (2006). Decision makers raise the inventory position of each item i to its order-up-to position S_i whenever a total of Q demand accumulated or T time units have elapsed, whichever occurs first. This policy is a hybrid of the continuous review (Q, S) policy, proposed by Renberg and Planche (1967), and the periodic review (R, T) policy, proposed by Atkins and Iyogun (1988). Thus, it features benefits of two separate policies. The comprehensive numerical study indicates that the proposed policy dominates the $P(s, S)$ policy, (Q, S) policy, $Q(s, S)$ policy, and (s, c, S) policy in 100 of 139 instances.

(R, S) policy. This policy is proposed by Bookbinder and Tan (1988) for controlling single-item inventory system. The time intervals between two consecutive orders R , and order-up-to positions S in replenishment periods are required

to be fixed at the beginning of the planning horizon. Decision makers raise the inventory position to S at the beginning of each replenishment period. This policy has been widely studied in the stream of literature on single-item lot-sizing problems. (Tarim and Kingsman, 2004, 2006) formulated a mixed integer programming (MIP) model for computing optimal (R, S) policy parameters. Tarim et al. (2011) relaxed the MIP model, and solved it as a shortest path problem which does not require the use of any MIP or Constraint Programming (CP) commercial solver. In addition, Özen et al. (2012) proposed a DP-based algorithm for solving small-size problems, and an approximation heuristic and a relaxation heuristic for tackling larger-size problems; Tunc et al. (2014) suggested a deterministic equivalent MIP model. Recently, Rossi et al. (2015) generalised the discussions above and developed a unified MILP model for approximating (R, S) policies by adopting the piecewise linear approximation technique in Rossi et al. (2014). Although various efficient modelling methods for computing (R, S) policy parameters were proposed, they generally control the single-item inventory system. The main purpose of this work is to apply the (R, S) policy for the multi-item inventory system. In the context of the JRP, a periodic-review (R, S) policy is adopted for each item.

The stochastic JRP is an open research area for the development of more efficient computational methods and control policies. In this study, we apply the periodic review (R, S) policy, originally proposed by Bookbinder and Tan (1988) for tackling single-item lot sizing problems, to JRPs with stochastic demand and fixed lead time. In the context of the JRP system, a periodic review (R, S) policy is adopted for each item. Note that when the demand is stationary stochastic, the (R, S) policy is the same as the MP policy proposed by Atkins and Iyogun (1988), which every T_n periods, raises the inventory position of item n to the order-up-to position R_n . However, the (R, S) policy also deals with non-stationary stochastic demand which was not addressed in Atkins and Iyogun (1988). In this paper, we present an MILP approach for approximating (R, S) policies under non-stationary stochastic demand. Nonlinear costs are approximated by leveraging the technique introduced in Rossi et al. (2014). Numerical experiments investigate the effectiveness of our approach against competing policies from the literature.

4.3 Problem description

Consider a periodic-review N -item inventory management system over a T -period planning horizon. We assume that demand d_t^n of item n , $n = 1, \dots, N$, in pe-

riod t , $t = 1, \dots, T$ are independently distributed random variables with known probability density function $g_t^n(\cdot)$.

We assume that ordering decisions are made at the beginning of each time period. There is a group fixed ordering cost K and an item-specific fixed ordering cost k^n . The group fixed ordering cost is incurred whenever an order is placed at a given time period, no matter which and how many items are included in this order. The item-specific fixed ordering cost is incurred whenever an order for item n is placed at a given time period, no matter how many items are included in this order.

We define Q_t^n as the quantity of item n ordered in period t , which will be received after lead time L^n . Then, the ordering cost of item n in period t with ordering quantity Q_t^n can be written as,

$$c_t^n(Q_t^n) = \begin{cases} k^n, & Q_t^n > 0, \\ 0, & Q_t^n = 0. \end{cases} \quad (4.1)$$

Let $c_t(\vec{Q}_t)$ denote the ordering cost of period t with ordering quantity vector $\vec{Q}_t = (Q_t^1, \dots, Q_t^N)$. $c_t(\vec{Q}_t)$ has the following structure

$$c_t(\vec{Q}_t) = \begin{cases} K + \sum_{n=1}^N c_t^n(Q_t^n), & \exists Q_t^n | Q_t^n > 0, \\ 0, & \text{otherwise.} \end{cases} \quad (4.2)$$

A penalty cost b^n is incurred for each unit of backorder demand for item n per period, and a holding cost h^n is charged for each unit of item n carried from one period to the next. The immediate penalty and holding cost of period t can be expressed as

$$L_t(\vec{y}) = \sum_{t=1}^n \left(b^n \cdot \mathbb{E}[\max(d_t^n - y^n, 0)] + h^n \cdot \mathbb{E}[\max(y^n - d_t^n, 0)] \right), \quad (4.3)$$

where vector $\vec{y} = (y^1, \dots, y^N)$ is the inventory level immediately after orders are received at the beginning of period t , and ‘‘E’’ denotes the expectation taken with respect to the random demand.

Let I_t^n denote the net inventory level of item n at the end of period t , which is also the opening inventory level of period $t + 1$, and $C_t(\vec{I}_{t-1})$ denote the expected total cost of an optimal policy over period t, \dots, T , given opening inventory level $\vec{I}_{t-1} = (I_{t-1}^1, \dots, I_{t-1}^N)$ at the beginning of period t . Note that there is no outstanding order at the beginning of the planning horizon. Then, $C_t(\vec{I}_{t-1})$ can be

written as,

$$C_t(\vec{I}_{t-1}) = \begin{cases} \min_{\vec{Q}_t} \{c_t(\vec{Q}_t) + L_t(\vec{I}_{t-1} + \vec{Q}_{t-\vec{L}}) + E[C_{t+1}(\vec{I}_{t-1} + \vec{Q}_{t-\vec{L}} - \vec{D}_t)]\}, & t \geq \vec{L} + 1, \\ \min_{\vec{Q}_t} \{c_t(\vec{Q}_t) + L_t(\vec{I}_{t-1}) + E[C_{t+1}(\vec{I}_{t-1} - \vec{D}_t)]\}, & \text{otherwise;} \end{cases} \quad (4.4)$$

where $\vec{D}_t = (d_t^1, \dots, d_t^N)$, $\vec{L} = (L^1, \dots, L^N)$, and

$$C_T(\vec{I}_{T-1}) = \begin{cases} \min_{\vec{Q}_T} \{c_T(\vec{Q}_T) + L_T(\vec{I}_{T-1} + \vec{Q}_{T-\vec{L}})\}, & t \geq \vec{L} + 1, \\ \min_{\vec{Q}_T} \{c_T(\vec{Q}_T) + L_T(\vec{I}_{T-1})\}, & \text{otherwise;} \end{cases} \quad (4.5)$$

represents the boundary condition. Moreover, let us define, $t = L^n + 1, \dots, T$,

$$G_t(\vec{I}_{t-1}) = L_t(\vec{I}_{t-1} + \vec{Q}_{t-\vec{L}}) + E[C_{t+1}(\vec{I}_{t-1} + \vec{Q}_{t-\vec{L}} - \vec{D}_t)]. \quad (4.6)$$

Example. We consider an instance in which the group fixed ordering cost $K = 10$, the item-specific ordering cost $k = 0$, the holding cost $h = 1$, the stock-out penalty cost $b = 5$. We control inventory for two items over a planning horizon of $T = 4$ periods. We assume that the demand of item n in period t follows a Poisson distribution with rate λ_t^n ; where $\lambda_t^1 = \lambda_t^2 = 3, 6, 9, 6$. For simplicity, we assume that the lead time is 0 for every item. The expected total cost, i.e. $C_1(\vec{I}_0)$, of an optimal policy, given initial inventory level $I_0^1 = I_0^2 = 0$, can be obtained via stochastic dynamic programming (SDP) and is equal to 65.4. In Fig. 4.1 we plot $G_1(\vec{I}_0)$ for $I_0^1 \in [0, 14]$ and $I_0^2 \in [0, 14]$.

4.4 An MILP model for approximating non-stationary stochastic (R, S) policies

In this section, we formulate the stochastic JRP problem under the (R, S) policy as an MILP model. Under the (R, S) policy, the replenishment periods and associated order-up-to positions are fixed at the beginning of the planning horizon, while actual order quantities are decided at the beginning of each replenish period. Note that in the context of JRP, a periodic-review (R, S) policy is adopted for each item. We first introduce a stochastic programming formulation in Section 4.4.1 and then we reformulate it as an MILP model in Section 4.4.2.




Figure 4.1: Expected total cost, i.e. $G_1(\vec{I}_0)$, contour plot for the two-item joint replenishment numerical example

4.4.1 A stochastic program

Consider the periodic-review N -item T -period JRP described in Section 4.3. We introduce binary variables δ_t and y_t^n , $t = 1, \dots, T$, and $n = 1, \dots, N$; δ_t takes value 1 if a group order is made in period t no matter how many types of items are involved, otherwise 0; y_t^n is set to 1 if item n is replenished in period t .

We further assume that the system is forced to place an order in period 1, and all orders should be received by the end of the planning horizon.

We reformulate the stochastic dynamic programming model in Section 4.3 as the stochastic program in Fig. 4.2.

The objective is to find the optimal replenishment plan so as to minimise the expected ordering costs, penalty costs, and holding costs of N items over the T -period planning horizon. Constraints (4.8) imply that if at least one item is ordered, then a group replenishment is issued. Constraints (4.9) force the system to replenish every item in period 1. Constraints (4.10) are inventory conservation constraints in periods $1, \dots, L^n$: inventory level at the end of period t is equal to the initial inventory level, minus demand realised up to period t . Constraints (4.11) ensure all replenishments are received by the end of the planning horizon. Constraints (4.12) are the inventory conservation constraints in periods $1 + L^n, \dots, T$: inventory level at the end of period t is equal to the initial inventory level, plus all orders received before the end of period t , minus demand realised up to period t . Constraints (4.13)- (4.16) state domains of y_t^n ,

$$\min \sum_{t=1}^T \left(K \cdot \delta_t + \sum_{n=1}^N (k^n \cdot y_t^n + b^n \mathbb{E}[\max(-I_t^n, 0)] + h^n \mathbb{E}[\max(I_t^n, 0)]) \right) \quad (4.7)$$

Subject to, $n = 1, \dots, N$,

$$\delta_t \geq y_t^n \quad t = 1, \dots, T \quad (4.8)$$

$$y_1^n = 1 \quad (4.9)$$

$$I_t^n = I_0^n - \sum_{j=1}^t d_j^n \quad t = 1, \dots, L^n \quad (4.10)$$

$$y_t^n = 0 \quad t = T - L^n, \dots, T \quad (4.11)$$

$$I_t^n = I_0^n + \sum_{i=1}^{t-L^n} Q_i^n - \sum_{j=1}^t d_j^n \quad t = L^n + 1, \dots, T \quad (4.12)$$

$$y_t^n = \begin{cases} 1, & Q_t^n > 0, \\ 0, & Q_t^n = 0. \end{cases} \quad (4.13)$$

$$Q_t^n \geq 0 \quad (4.14)$$

$$\delta_t = \{0, 1\} \quad (4.15)$$

$$I_t^n \in \mathcal{R} \quad (4.16)$$

Figure 4.2: Stochastic programming formulation of the JRP.

Q_t^n , δ_t , and I_t^n .

4.4.2 An MILP model

The stochastic programming formulation in Fig. 4.2 can be reformulated into an MILP model via the piecewise approximation approach in Rossi et al. (2014).

In the rest of this paper, let “ \sim ” denote the expectation operator. We introduce the first order loss function

$$\mathcal{L}(x, \omega) = \int_x^\infty \max(t - x, 0) g_\omega(t) d(t)$$

and its complementary function

$$\hat{\mathcal{L}}(x, \omega) = \int_{-\infty}^x \max(x - t, 0) g_\omega(t) d(t),$$

where ω is a random variable with probability density function $g_\omega(\cdot)$, and x is a scalar variable.

Consider a partition of the support Ω of ω into W disjoint subregions $\Omega_1, \dots, \Omega_W$, the probability mass $p_i = \Pr\{\omega \in \Omega_i\} = \int_{\Omega_i} g_\omega(t) d(t)$, and the conditional ex-

pectation $E[d_{jt}|\Omega_i] = \frac{1}{p_i} \int_{\Omega_i} t g_{\omega}(t) dt$, $i = 1, \dots, W$. By applying Jensen's lower bound,¹ $L(x, \omega)$ and $\hat{L}(x, \omega)$ can be approximated as piecewise linear functions, as presented in the following lemma.

Lemma 4.4.1. *For the first order loss function and its complementary function, the lower bounds \mathcal{L}_{lb} and $\hat{\mathcal{L}}_{lb}$, where $E[\omega|\Omega_i] \leq x \leq E[\omega|\Omega_{i+1}]$, $i = 1, \dots, W$,*

$$\mathcal{L}(x, \omega) \geq \mathcal{L}_{lb}(x, \omega) = x \sum_{k=1}^i p_k + \sum_{k=1}^i p_k E[\omega|\Omega_k] + (x - \tilde{\omega}), \quad (4.17)$$

$$\hat{\mathcal{L}}(x, \omega) \geq \hat{\mathcal{L}}_{lb}(x, \omega) = x \sum_{k=1}^i p_k + \sum_{k=1}^i p_k E[\omega|\Omega_k] \quad (4.18)$$

are piecewise linear functions with $W + 1$ segments.

We introduce two sets of variables $\tilde{B}_t^n \geq 0$ and $\tilde{H}_t^n \geq 0$ represent lower bounds of $E[\max(-I_t^n, 0)]$ and $E[\max(I_t^n, 0)]$, $t = 1, \dots, T$, $n = 1, \dots, N$. Then, the objective function (4.7) in Fig. 4.2 can be rewritten as

$$\min \sum_{t=1}^T \left(K \cdot \delta_t + \sum_{n=1}^N (k^n \cdot y_t^n + b^n \tilde{B}_t^n + h^n \tilde{H}_t^n) \right). \quad (4.19)$$

We next construct constraints by separating the discussion into two parts.

The first part involves periods $1, \dots, L^n$, $n = 1, \dots, N$, where no order is received. Recall that there is no outstanding order at the beginning of the planning horizon, and the system is forced to issue an order in period 1, then the inventory level I_t^n must equal to the initial inventory level of item n at the beginning of the planning horizon, minus the demand convolution over periods $1, \dots, t$, i.e., $I_t^n = I_0^n - d_{1,t}^n$, where $d_{1,t}^n$ is the demand convolution of item n over periods $1, \dots, t$, i.e., $d_{1,t}^n = d_1^n + \dots + d_t^n$. We rewrite the expected back-orders and excess on-hand stocks using the first order loss function and its complementary function, $\mathcal{L}(I_0^n, d_{1,t}^n)$ and $\hat{\mathcal{L}}(I_0^n, d_{1,t}^n)$. By applying Lemma 4.4.1, \tilde{B}_t^n and \tilde{H}_t^n can be written as follows, $t = 1, \dots, L^n$, $n = 1, \dots, N$, $i = 1, \dots, W$,

$$\tilde{B}_t^n \geq -\tilde{I}_t^n + \sum_{k=1}^i p_k I_0^n - \sum_{k=1}^i p_k E[d_{1,t}^n|\Omega_i], \quad (4.20)$$

$$\tilde{H}_t^n \geq \sum_{k=1}^i p_k I_0^n - \sum_{k=1}^i p_k E[d_{1,t}^n|\Omega_i]. \quad (4.21)$$

¹Similarly, the Edmundson-Madansky upper bound can be applied for approximating the expected excess inventory and back-orders as well, for further details refer to (Rossi et al., 2014).

Additionally, constraints (4.10) in Fig. 4.2 can be rewritten as,

$$\tilde{I}_t^n + \tilde{d}_t^n - \tilde{I}_{t-1}^n = 0, \quad t = 1, \dots, L^n. \quad (4.22)$$

The second part involves periods $1 + L^n, \dots, T$, $n = 1, \dots, N$. Consider a single cycle of item n over periods i, \dots, j , in which a single order is received at the beginning of period i , and the next order will be received at the beginning of period $j + 1$. Since the lead time of item n is L^n , the order that arrives in period i must be issued in period $i - L^n$ with order-up-to position $S_{i-L^n}^n$. Thus, I_t^n , $t = \{i, \dots, j\}$, must equal to the order-up-to position $S_{i-L^n}^n$, minus the demand convolution over periods $i - L^n, \dots, t$, i.e. $I_t^n = S_{i-L^n}^n - d_{i-L^n,t}^n$.

We introduce a binary variable P_{jt}^n which is set to one if the most recent order received before period t arrived in period j , where $j \leq t$, $j = 1 + L^n, \dots, t$, $t = 1 + L^n, \dots, T$, and $n = 1, \dots, N$; and we introduce the following constraints, $t = 1 + L^n, \dots, T$, $n = 1, \dots, N$,

$$\sum_{j=1+L^n}^t P_{jt}^n = 1, \quad (4.23)$$

$$P_{j,t}^n \geq y_{j-L^n}^n - \sum_{k=j-L^n+1}^{t-L^n} y_k^n, \quad j = 1 + L^n, \dots, t. \quad (4.24)$$

Constraints (4.23) indicate that the most recent order received before period t arrived in period j , $1 + L^n \leq j \leq t$. Constraints (4.24) identify uniquely the period in which the most recent order received before period t has been received. Therefore, the inventory level $I_t^n = \sum_{j=1+L^n}^t (S_{j-L^n}^n - d_{j-L^n,t}^n) P_{jt}^n$, where $t = 1 + L^n, \dots, T$, and $S_{j-L^n}^n$ represents the order-up-to position of item n in period $j - L^n$. We write the back-orders and excess inventory as the first order loss function and its complementary, $\sum_{j=1+L^n}^t \mathcal{L}(S_{j-L^n}^n, d_{j-L^n,t}^n) P_{jt}^n$ and $\sum_{j=1+L^n}^t \hat{\mathcal{L}}(S_{j-L^n}^n, d_{j-L^n,t}^n) P_{jt}^n$. By applying Lemma 4.4.1, \tilde{B}_t^n and \tilde{H}_t^n can be written as, $t = 1 + L^n, \dots, T$, $n = 1, \dots, N$, $i = 1, \dots, W$,

$$\tilde{B}_t^n \geq -\tilde{I}_t^n + (\tilde{I}_t^n + \sum_{j=1+L^n}^t \tilde{d}_{j-L^n,t}^n P_{jt}^n) \sum_{k=1}^i p_k - \sum_{j=1+L^n}^t \sum_{k=1}^i p_k E[d_{j-L^n,t}^n | \Omega_i] P_{jt}^n, \quad (4.25)$$

$$\tilde{H}_t^n \geq (\tilde{I}_t^n + \sum_{j=1+L^n}^t \tilde{d}_{j-L^n,t}^n P_{jt}^n) \sum_{k=1}^i p_k - \sum_{j=1+L^n}^t \sum_{k=1}^i p_k E[d_{j-L^n,t}^n | \Omega_i] P_{jt}^n. \quad (4.26)$$

Note that $S_{j-L^n}^n = \tilde{I}_t^n + \tilde{d}_{j-L^n,t}^n$. In addition, constraints (4.12)-(4.14) in Fig. 4.2

can be reformulated as follows,

$$y_{t-L^n}^n = 0 \rightarrow \tilde{I}_t^n + \tilde{d}_t^n - \tilde{I}_{t-1}^n = 0, \quad t = 1 + L^n, \dots, T, \quad (4.27)$$

$$\tilde{I}_t^n + \tilde{d}_t^n - \tilde{I}_{t-1}^n \geq 0, \quad t = 1 + L^n, \dots, T. \quad (4.28)$$

We now present the overall model in Fig. 4.3. The objective function (4.29) minimise the expected group fixed ordering costs, item-specific fixed ordering costs, penalty costs, and holding costs of N items over the T -period planning horizon. Constraints (4.30) imply an individual item can only be included in a group replenishment if that replenishment is made. Constraints (4.31) - (4.32) assume that the first order is issued at the beginning of period 1, and there is no outstanding replenishment at the beginning of the planning horizon. Constraints (4.33) - (4.34) represent the expected back-orders and on-hand stocks of item n over periods $1, \dots, L^n$. Constraints (4.35) state all orders are received by the end of the planning horizon. Constraints (4.36) - (4.37) are inventory balance constraints. Constraints (4.38) - (4.39) ensure the most recent replenishment that has arrived before period t was received in period j . Constraints (4.40) - (4.41) represent the expected back-orders and on-hand stocks of item n over periods $1 + L^n, \dots, T$. Constraints (4.42) - (4.44) indicate domains of binary variables δ_t^n , y_t^n , and P_{jt}^n .

By solving the model in Fig. 4.3, the optimal replenishment plan including group replenish periods δ_t , and item-specific replenish periods y_t^n , and the item-specific order-up-to positions $S_t^n = \tilde{I}_{t+L^n}^n + \tilde{d}_{t,t+L^n}^n$ are obtained, for $t = 1, \dots, T$, and $n = 1, \dots, N$.

Example. We demonstrate the modelling strategy behind the MILP model on a 5-item 10-period example. It is assumed that the demand is Poisson-distributed with rate λ_t^n presented in Table 4.1. The initial inventory level is taken as zero. Other parameters are: $K = 500$, $b = 10$, $h = 2$, $k^n = 120, 100, 80, 120, 150$, and $L^n = 1, 2, 3, 1, 3$. We employ eleven segments in the piecewise-linear approximations of \tilde{B}_t^n and \tilde{H}_t^n , for $n = 1, \dots, 5$, and $t = 1, \dots, 10$.

The resulting expected total cost is 14236.24. Replenishment plans of each item are presented in Fig. 4.4. Items 1, 2 and 4 are replenished in periods 1, 3, 5, and 8; while items 3 and 5 are replenished only in periods 1, 3, and 5 since orders in period 8 could not be received by the end of the planning horizon. Additionally, item 1 is expected to be ordered every two periods with the same order-up-to position 123 by the nature of stationary demand, while it is ordered

$$\min \sum_{t=1}^T \left(K \cdot \delta_t + \sum_{n=1}^N (k^n \cdot y_t^n + h^n \tilde{H}_t^n + b^n \tilde{B}_t^n) \right) \quad (4.29)$$

Subject to, $n = 1, \dots, N$

$$\delta_t \geq y_t^n \quad t = 1, \dots, T \quad (4.30)$$

$$y_1^n = 1 \quad (4.31)$$

$$\tilde{I}_t^n + \tilde{d}_t^n - \tilde{I}_{t-1}^n = 0 \quad t = 1, \dots, L^n \quad (4.32)$$

$$B_t^n \geq -\tilde{I}_t^n + \sum_{k=1}^i p_k I_0^n - \sum_{k=1}^i p_k E[d_{1,t}^n | \Omega_i], \quad t = 1, \dots, L^n, i = 1, \dots, W \quad (4.33)$$

$$H_t^n \geq \sum_{k=1}^i p_k I_0^n - \sum_{k=1}^i p_k E[d_{1,t}^n | \Omega_i], \quad t = 1, \dots, L^n, i = 1, \dots, W \quad (4.34)$$

$$y_t^n = 0 \quad t = T - L^n, \dots, T \quad (4.35)$$

$$\tilde{I}_t^n + \tilde{d}_t^n - \tilde{I}_{t-1}^n \geq 0 \quad t = 1 + L^n, \dots, T \quad (4.36)$$

$$y_{t-L^n}^n = 0 \rightarrow \tilde{I}_t^n + \tilde{d}_t^n - \tilde{I}_{t-1}^n = 0 \quad t = 1 + L^n, \dots, T \quad (4.37)$$

$$\sum_{j=1+L^n}^t P_{jt}^n = 1 \quad t = 1 + L^n, \dots, T \quad (4.38)$$

$$P_{j,t}^n \geq y_{j-L^n}^n - \sum_{k=j-L^n+1}^t y_k^n \quad t = 1 + L^n, \dots, T, j = 1, \dots, t \quad (4.39)$$

$$B_t^n \geq -\tilde{I}_t^n + (\tilde{I}_t^n + \sum_{j=1+L^n}^t \tilde{d}_{j-L^n,t}^n P_{jt}^n) \sum_{k=1}^i p_k - \sum_{j=1+L^n}^t \sum_{k=1}^i p_k E[d_{j-L^n,t}^n | \Omega_i] P_{jt}^n \quad t = 1 + L^n, \dots, T, i = 1, \dots, W \quad (4.40)$$

$$H_t^n \geq (\tilde{I}_t^n + \sum_{j=1+L^n}^t \tilde{d}_{j-L^n,t}^n P_{jt}^n) \sum_{k=1}^i p_k - \sum_{j=1+L^n}^t \sum_{k=1}^i p_k E[d_{j-L^n,t}^n | \Omega_i] P_{jt}^n \quad t = 1 + L^n, \dots, T, i = 1, \dots, W \quad (4.41)$$

$$\delta_t = \{0, 1\} \quad t = 1, \dots, T \quad (4.42)$$

$$y_t^n = \{0, 1\} \quad t = 1, \dots, T \quad (4.43)$$

$$P_{jt}^n = \{0, 1\} \quad t = 1 + L^n, \dots, T, j = 1 + L^n, \dots, t \quad (4.44)$$

Figure 4.3: MILP model for approximating (R, S) policies

up to a higher position 164 in period 5 to cover demand in the next 3 periods in order to coordinate with other items.




Figure 4.4: Replenishment plans of the 5-item 10-period example

λ_t^n \ period	1	2	3	4	5	6	7	8	9	10
item 1	40	40	40	40	40	40	40	40	40	40
2	5	64	29	54	70	50	54	45	13	50
3	40	55	72	86	78	51	42	38	30	26
4	41	58	75	63	40	35	33	18	29	39
5	45	40	22	31	38	46	59	62	46	40

Table 4.1: Demand rates λ_t^n of the 5-item 10-period example

4.5 MILP model for approximating the optimal (σ, \vec{S}) policies

Since Scarf (1960) proved the optimality of (s, S) policy for the classical single-item inventory problem, there have been several attempts to prove the optimality of (σ, \vec{S}) policy for the multi-item inventory problem, e.g.: (Johnson, 1967; Kalin, 1980; Ohno and Ishigaki, 2001; Gallego and Sethi, 2005). However, the computation of optimal (σ, \vec{S}) policy parameters is still a difficult task. In this section we show how the MILP model proposed in Section 4.4.2 can be used to approximate the optimal replenishment plan under (σ, \vec{S}) policy for the JRP.

Definition 4.5.1 (Gallego and Sethi (2005)). Function $f(\cdot) : \mathcal{R}^N \rightarrow \mathcal{R}$ is K -convex if

$$f(ax + (1-a)z) \leq af(x) + (1-a)[f(z) + \mathbf{K}\delta(z-x)],$$

where $x \leq z$, $a \in [0, 1]$, and $\mathbf{K}\delta(z-x)$ is defined as follows,

$$\mathbf{K}\delta(z-x) = K\delta(e'x) + \sum_{n=1}^N k^n \delta(x_n),$$

where $e' = (1, 1, \dots, 1)' \in \mathcal{R}^N$, $\delta(0) = 0$, and $\delta(y) = 1$ for all $y > 0$.

Gallego and Sethi (2005) showed the optimal policy for the joint setup cost case by studying function

$$G_t(\vec{y}) = L_t(\vec{y}) + C_{t+1}(\vec{y} - \vec{d}_t). \quad (4.45)$$

Consider a continuous K -convex function $G_t(\cdot)$, and it has a global minimum at \vec{S}_t . Define set $\Sigma = \{\vec{I}_{t-1} \leq \vec{S}_t | G_t(\vec{I}_{t-1}) \leq G_t(\vec{S}_t) + K\}$, and set $\sigma = \{\vec{I}_{t-1} \leq$

$\vec{S}_t | \vec{I}_{t-1} \notin \Sigma\}$. Then, the optimal replenish plan is to order up to \vec{S}_t if opening inventory levels $\vec{I}_{t-1} \in \sigma$ and $\vec{I}_{t-1} \leq \vec{S}_t$; otherwise, not to order (Gallego and Sethi, 2005).

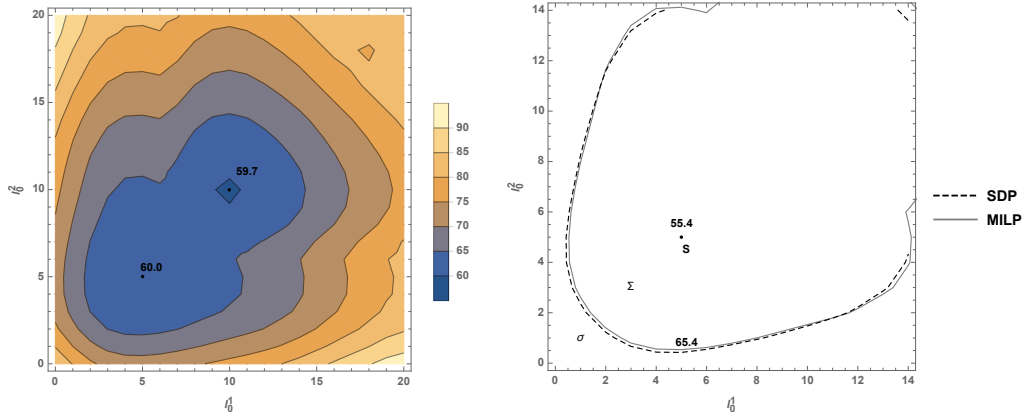
We next show that the MILP model in Fig. 4.3 can be adjusted to approximate set σ and \vec{S} . Since Gallego and Sethi (2005) showed that the (σ, \vec{S}) policy is optimal when expected total costs only consist of group fixed ordering costs, holding costs and penalty costs, we first drop the item-specific fixed ordering cost, i.e., $k^n \cdot y_t^n$, in the objective function (4.29). Additionally, since the lead time is not considered in Gallego and Sethi (2005), we then set the lead time of all items to 0, i.e.: $L^n = 0$, $n = 1, \dots, N$, and drop constraints (4.31) - (4.35).

Due to the complexity of σ , it is impractical to derive a closed form expression for it. Alternatively, one may propose a strategy to determine whether given initial inventory levels $\vec{I}_0 \in \sigma$. By solving our modified MILP model over planning horizon k, \dots, T , $k = 1, \dots, T$, we observe the minimised expected total cost $G_k(\vec{S}_k)$, order-up-to levels \vec{S}_k , and the order decision δ_k . If $\delta_k = 1$, then $\vec{I}_{k-1} \in \sigma$, which means decision makers have to order up to \vec{S}_t ; otherwise, $\vec{I}_{k-1} \in \Sigma$, which means decision makers do not need to place orders. Therefore, our MILP model can be used to determine whether decision makers need to place orders with given initial inventory levels. Moreover, by repeating this procedure, one can approximate the optimal replenishment strategy for every period $k = 1, \dots, T$.

Example. We illustrate the concept introduced on the 2-item 4-period example presented in Section 4.3. Assuming the initial inventory level $\vec{I}_0^1 \in [0, \dots, 20]$, and $\vec{I}_0^2 \in [0, \dots, 20]$, we plot the expected total cost contours, obtained via the modified MILP in Fig. 4.5(a). Note that there are two similar minima, which is expected since the ordering cost is relatively small and the demand variance is large. We plot set σ and \vec{S} obtained via the modified MILP model, and compare them with that obtained via stochastic dynamic programming in Fig. 4.5(b). The optimal policy is to place an order whenever inventory levels $\vec{I}_0 = (I_0^1, I_0^2)$ fall in set σ , and not to place an order if \vec{I}_0 fall in Σ . We observe that set σ and \vec{S} obtained via the modified MILP model neatly approximate those obtained via stochastic dynamic programming.

4.6 Computational Experiments

In this section we assess the cost performance of the (R, S) policy by comparing its cost performance against (Q, S, T) policy (Özkaya et al., 2006), $Q(s, S)$ policy (Nielsen and Larsen, 2005), $P(s, S)$ policy (Viswanathan, 1997), (Q, S) policy



(a) Expected total cost contour plot obtained via MILP approximation (b) Plot of expected total costs obtained via MILP and SDP

Figure 4.5: Plot of expected total costs for the two-item joint replenishment numerical example

(Pantumsinchai, 1992), MP policy (Atkins and Iyogun, 1988), $(s, c, S)_M$ policy (Melchioris, 2002), and $(s, c, S)_F$ policy (Federgruen et al., 1984), on data sets of Atkins and Iyogun (1988) and Viswanathan (1997). These data sets consider stationary demand over an infinite horizon. Unfortunately, computing (R, S) policy parameters for infinite horizon JRPs via our MILP model is computationally expensive; however, since demand is stationary, it is possible to derive an efficient shortest path reformulation, which we present in 4.A and we use in our computational study.

Computational experiments are conducted by using IBM ILOG CPLEX Optimization Studio 12.7 and Matlab R2016a on a 3.20 GHz Intel Core i5-6500 CPU with 16.0 GB RAM, 64 bit machine.

Since the shortest path reformulation operates over a finite horizon, in order to compare the cost performance of the (R, S) policy with continuous-review (s, c, S) , (Q, S) , and (Q, S, T) policy, we discretize each time period into 20 small periods. We consider a planning horizon length of 6.6 periods for a total of 132 small periods. For each test instance, we first obtain the optimal replenishment plan by solving the shortest path reformulation presented in 4.A. The computational time is limited to 5 minutes, if a timeout occurs, the best solution available is adopted. Next, we simulate the expected average cost of each test instance via Monte Carlo Simulation (100,000 replications). Finally, we compare the average cost per small period against the average cost under existing policies.

The data set of Atkins and Iyogun (1988) assumes that the demand of each item follows stationary Poisson distribution with rate λ^n , $n = 1, \dots, 12$. The item-specific fixed ordering cost K^n , expected demand λ^n , and lead time L^n are

displayed in Table 4.2. Items share the same penalty cost $b = 30$, holding cost $h \in \{2, 6, 20\}$, and group fixed ordering cost $K \in \{20, 50, 100, 150, 500\}$.

items	1	2	3	4	5	6	7	8	9	10	11	12
K^n	10	10	20	20	40	20	40	40	60	60	80	80
λ^n	40	35	40	40	40	20	20	20	28	20	20	20
L^n	0.2	0.5	0.2	0.1	0.2	1.5	1.0	1.0	1.0	1.0	1.0	1.0

Table 4.2: K^n , λ^n , and L^n of data set Atkins and Iyogun (1988)

The data set of Atkins and Iyogun (1988) contains some unusual lot sizing instances; more specifically, instances for which the group as well as item fixed ordering costs become negligible in comparison to holding costs. In the lot-sizing literature the fixed ordering cost is commonly assumed to be greater than the holding cost (see Axsäter, 2010, p. 62, Property 2); moreover, the penalty cost should not be smaller than the holding cost. Additionally, we observe that the fixed ordering cost should be no greater than the penalty cost, otherwise the inventory system tends to place orders in every period instead of penalising backorders. To focus on meaningful lot sizing instances — instances in which a trade off between fixed ordering and holding/penalty cost is sought — we filter test instances of the data set of Atkins and Iyogun (1988) by using the following conditions: $K > b \geq h$. We also check the order frequency in each period and we discard instances in which orders are issued too frequently — i.e. instance in which a replenishment is issued more than twice per time period, as it turns out that for these instances order coordination is straightforward due to negligible item fixed ordering costs: if a group order is placed, all items are ordered. We present the filtered computational results in Table 4.3, and complete results in the Appendix (Table 4.5).

K	b	h	(R, S)	Average cost improvement $\Delta\%$							
				(Q, S, T)	$Q(s, S)$	$P(s, S)$	(Q, S)	MP	$(s, c, S)_M$	$(s, c, S)_F$	
50	30	2	936.94	-0.91	-0.84	-0.33	4.38	0.68		0.79	2.14
100	30	2	990.50	-0.05	-0.45	0.75	2.57	1.77		4.39	6.81
150	30	2	1046.56	-0.24	-1.01	-0.35	0.52	0.65		5.68	8.36
200	30	2	1072.97	1.32	0.47	1.11	1.34	2.12		8.34	12.31
100	30	6	1639.75	-0.23	-1.52	-1.02	2.15	0.00		1.24	3.31
150	30	6	1707.05	0.64	-0.60	-0.07	1.46	0.95		2.34	6.68
200	30	6	1766.38	1.16	0.08	0.65	1.17	1.67		3.08	9.04
150	30	20	2718.47	0.77	4.32	-1.26	1.27	-0.21		-0.59	6.20
200	30	20	2812.52	-3.23	0.14	-0.72	0.77	0.34		0.25	8.34
Average cost improvement $\Delta\%$				-0.09	0.07	-0.14	1.74	0.89		2.84	7.02

Table 4.3: Computational results on the data set of Atkins and Iyogun (1988)

Let $\Delta\%$ denote the percentage gap between the expected average cost of existing policies and that of the proposed (R, S) policy, over the expected average

cost of the (R, S) policy. By definition, a positive $\Delta\%$ represents the (R, S) policy outperforms existing policy. Note that expected average costs under (Q, S, T) , $Q(s, S)$, $P(s, S)$, (Q, S) , and $(s, c, S)_M$ policy are obtained from Özkaya et al. (2006), that of $(s, c, S)_F$ policy is obtained from Melchioris (2002), and that of MP policy is obtained from Viswanathan (1997).

We observe that the (R, S) policy fully dominates all policies in 2 of 9 test instances; (Q, S, T) is the best policy in 2 instances; $Q(s, S)$ is the best policy in 4 instances; $P(s, S)$ is the best policy in 1 instance. Moreover, the (R, S) policy outperforms the (Q, S) and $(s, c, S)_F$ policy, and there is no dominant policy, on all test instances. The average cost improvement $\Delta\%$ increases with the increase of group fixed ordering cost, and decreases with the increase of holding cost compared with $(s, c, S)_M$ and $(s, c, S)_F$ policy. That means an increase in group fixed ordering cost or a decrease in holding cost improves the cost performance of (R, S) policy. It is difficult to make a general remark with respect to group fixed ordering cost and holding cost compared with (Q, S, T) , $Q(s, S)$, $P(s, S)$, (Q, S) , and MP policy. On average, the (R, S) policy performs better than $Q(s, S)$, (Q, S) , MP , $(s, c, S)_M$, and $(s, c, S)_F$ policy with an average improvement of 0.07%, 1.74%, 0.89%, 2.84%, and 7.02%, respectively; however, the (Q, S, T) and $P(s, S)$ policies performs slightly better than the (R, S) policy with an average improvement of 0.09% and 0.14%, respectively.

Viswanathan (1997) adopts the same experimental setup as Atkins and Iyogun (1988), except $h \in \{2, 6, 10, 200, 600, 1000\}$, $K \in \{20, 50, 100, 200, 500\}$, and $b \in \{10, 50, 100, 200, 1000, 5000, 10000, 20000\}$.

We filter the computational results by using the same conditions previously adopted. We present computational results of the (R, S) policy on the data set of Viswanathan (1997) in Table 4.4, and complete computational results in Table 4.6. We observe that the (R, S) policy dominates 13 of 31 test instances; (Q, S, T) is the best policy in 13 instances; $Q(s, S)$ is the best policy in 9 instances; $P(s, S)$ is the best policy in 1 instances. There is once more no dominant policy on all test instances. Regarding the comparison with other policies, the average cost improvement $\Delta\%$ decreases as the penalty cost increases; while there is no obvious trend with respect to the group fixed ordering cost, and penalty cost. On average, the (R, S) policy performs better than $Q(s, S)$, $P(s, S)$, (Q, S) , MP , and $(s, c, S)_F$ policy with average cost improvements of 0.37%, 0.37%, 1.81%, 1.41%, and 1.67%; while the (Q, S, T) policy performs slightly better than the (R, S) policy with average cost improvement 0.19%.

Even though the (R, S) policy does not fully dominate other competing poli-

K	b	h	(R, S)	Average cost improvement $\Delta\%$					
				(Q, S, T)	$Q(s, S)$	$P(s, S)$	(Q, S)	MP	$(s, c, S)_F$
20	10	2	772.25	-0.03	0.48	0.76	8.30	1.79	1.80
50	10	2	813.94	-0.48	0.12	0.62	0.47	1.64	1.74
100	10	2	861.05	0.23	0.70	1.17	3.68	2.20	2.38
200	10	2	932.86	1.62	1.83	2.38	2.88	3.42	3.73
500	10	2	1131.42	0.14	0.14	0.59	0.18	1.60	2.12
20	10	6	1166.06	0.85	2.84	0.01	7.99	1.08	1.04
50	10	6	1222.82	-0.15	1.83	0.62	5.53	1.68	1.73
100	10	6	1283.92	1.33	2.50	1.26	4.49	2.34	2.46
200	10	6	1413.72	0.30	1.23	1.02	1.82	2.10	2.33
500	10	6	1658.48	2.26	2.20	2.52	2.30	3.59	4.03
50	10	10	1420.63	1.57	5.30	-0.03	5.88	1.07	1.07
100	10	10	1497.96	1.67	4.28	0.75	4.37	1.87	1.93
200	10	10	1637.27	0.66	2.18	1.15	2.16	2.28	2.44
500	10	10	1935.07	1.60	1.60	1.79	1.60	2.90	3.27
100	50	2	1043.31	-1.95	-0.79	-0.23	1.98	0.78	0.92
200	50	2	1132.61	-1.29	-0.48	0.30	0.50	1.31	1.97
500	50	2	1327.95	0.08	0.08	0.82	0.13	1.83	2.30
100	50	6	1794.60	-1.37	-2.65	-2.09	0.94	-1.09	-0.97
200	50	6	1938.25	-0.27	-1.56	-0.89	-0.05	0.13	0.34
500	50	6	2244.01	-0.27	-0.27	0.43	-0.26	1.44	1.87
200	50	10	2448.79	-3.83	-2.11	-1.55	-0.75	-0.53	-0.34
500	50	10	2796.29	0.35	0.35	0.97	0.35	2.00	2.40
200	100	2	1200.38	-1.61	-0.94	-0.11	-0.01	0.90	1.13
500	100	2	1406.67	-0.76	-0.83	0.16	0.16	1.17	1.60
200	100	6	2106.78	0.44	-1.23	-0.48	0.94	0.54	0.73
500	100	6	2449.51	-0.88	-0.88	-0.07	-0.07	0.94	1.33
200	100	10	2728.08	-3.41	-1.90	-1.29	-0.49	-0.27	-0.10
500	100	10	3108.05	0.22	0.22	0.94	0.94	1.96	2.33
500	200	2	1470.29	-0.90	-0.90	0.05	0.05	1.05	1.45
500	200	6	2620.77	-0.91	-0.91	0.08	0.08	1.09	1.45
500	200	10	3421.28	-0.94	-0.94	-0.04	-0.04	0.97	1.30
Average cost improvement $\Delta\%$				-0.19	0.37	0.37	1.81	1.41	1.67

Table 4.4: Computational results on the data set of Viswanathan (1997)

cies, it presents a key advantage: *in contrast to all other policies in the literature, it is able to tackle stationary as well as nonstationary demand.*

4.7 Conclusion

In this paper, we presented a mathematical programming approach for controlling the multi-item inventory system with joint replenishment under the (R, S) policy. We first present an MILP-based model for approximating optimal (R, S) policies, which is built upon the piecewise-linear approximation technique proposed by Rossi et al. (2014). We further demonstrate that the MILP model can be used to approximate the (σ, \vec{S}) policy.

We conducted an extensive computational study comprising 40 instances. We

first evaluated our approach on the data set of Atkins and Iyogun (1988). This evaluation demonstrates that the (R, S) policy fully dominates other competing policies in the literature in 2 out of 9 test instances considered. The (R, S) policy performs better than $Q(s, S)$, (Q, S) , MP , $(s, c, S)_M$, and $(s, c, S)_F$ policies with an average improvement of 0.07%, 1.74%, 0.89%, 2.84%, and 7.02%, respectively; however, the (Q, S, T) and $P(s, S)$ policies performs slightly better than the (R, S) policy with an average improvement of 0.09% and 0.14%. Computational experiments on the data set of Viswanathan (1997) indicates that (R, S) is the best policy in 13 out of 31 test instances. (R, S) performs better than $Q(s, S)$, $P(s, S)$, (Q, S) , MP , and $(s, c, S)_F$ policies with average cost improvements of 0.37%, 0.37%, 1.81%, 1.41%, and 1.67%; while (Q, S, T) policy performs slightly better than it with an average cost improvement 0.19%. Most importantly, the (R, S) policy comes with the additional advantage of being able to tackle stationary and nonstationary demand. Future research may focus on investigating the cost performance of (R, S) policy in a rolling horizon setting.

Appendix

4.A Shortest path reformulation for approximating stationary stochastic (R, S) policies

In this section we present an efficient shortest path reformulation for computing stationary (R, S) policies.

Consider a network $\mathcal{G} = (\mathcal{N}, \mathcal{A})$ with nodes $\mathcal{N} = \{1, \dots, T\}$ representing time periods, and arc (i, j) between each pair of (i, j) representing a possible decision to issue an order in period i to satisfy demand in periods i, \dots, j . Assigning a cost to this arc, solving the optimisation problem in Fig. 4.4 is equivalent to finding the shortest path between nodes 1 and T in the network \mathcal{G} . In the rest of this section, we first present how to compute the cost of each arc, and then present the shortest path reformulation.

Consider a replenishment cycle i, \dots, j , where the only order is issued in period i with order-up-to position S_{ij}^n , and the next order is issued in period $j + 1$, for $i = 1, \dots, T$, $j = i, \dots, T$, $n = 1, \dots, N$. We assume d_t^n follows Poisson

distribution with rate λ^n . Then, S_{ij}^n is calculated by Askin (1981),

$$\sum_{t=i}^j G_{d_{i,t+L^n}^n}(S_{i,j}^n) = \frac{(j-i+1) \cdot b^n}{h^n + b^n}. \quad (4.46)$$

Note that the order-up-to position $S_{i,j}^n$ actually accounts for demand variances over periods $i, \dots, j + L^n$, which is reflected on the cumulative distribution function $G_{d_{i,t+L^n}^n}(\cdot)$ on the left-hand-side of Eq. (4.46).

Since the demand of item n follows Poisson distribution with rate λ^n , we could approximate the cost of the replenishment cycle i, \dots, j by that of the cycle $i + L^n, \dots, j + L^n$ as shown in Fig. 4.6. As a result, the cycle cost c_{ij}^n can be calculated as follows,

$$c_{ij}^n = k^n + h^n \sum_{t=i}^j \hat{\mathcal{L}}(S_{i,j}^n - L^n \lambda^n, d_{it}) + b^n \sum_{t=i}^j \mathcal{L}(S_{i,j}^n - L^n \lambda^n, d_{it}). \quad (4.47)$$




Figure 4.6: Expected inventory curve under (R, S) policy.

At the beginning of the planning horizon, the initial inventory level is I_0^n . We check the cost of not issuing an order in period 1, \bar{c}_{1j}^n , and update c_{1j}^n with \bar{c}_{1j}^n if $\bar{c}_{1j}^n \leq c_{1j}^n$, for $j = 1, \dots, T$.

$$\bar{c}_{1j}^n = h^n \cdot \sum_{t=1}^j \hat{\mathcal{L}}(I_0^n, d_{1t}) + b^n \cdot \sum_{t=1}^j \mathcal{L}(I_0^n, d_{1t}). \quad (4.48)$$

Additionally, we introduce an auxiliary binary variable P_j^n , which is equal to 1 if an order is placed in period 1 to satisfy demand in cycle $1, \dots, j$, otherwise 0.

We now present the shortest path reformulation in Fig. 4.7. Let binary variable Y_{ij}^n equal to 1 if an order is issued in period i to cover demand in periods

i, \dots, j , otherwise 0. The objective is to find the optimal replenishment plan that minimising the expected group fixed order costs, item-specific fixed order costs, holding costs and penalty costs over periods $1, \dots, T$ for items $1, \dots, N$.

$$\min \sum_{i=1}^T K \cdot \delta_i + \sum_{n=1}^N \sum_{i=1}^T \sum_{j=i}^T c_{ij}^n \cdot Y_{ij}^n \quad (4.49)$$

subject to, $n = 1, \dots, N$,

$$\delta_1 \geq \sum_{j=1}^T Y_{1j}^n \cdot P_j^n \quad (4.50)$$

$$\delta_i \geq \sum_{j=i}^T Y_{ij}^n \quad i = 2, \dots, T \quad (4.51)$$

$$\sum_{j=1}^T Y_{1j}^n = 1 \quad (4.52)$$

$$\sum_{j=i}^T Y_{ij}^n - \sum_{k=1}^{i-1} Y_{ki}^n = 0 \quad i = 2, \dots, T-1 \quad (4.53)$$

$$\sum_{i=1}^T Y_{iT}^n = 1 \quad (4.54)$$

Figure 4.7: Shortest path formulation for approximating stationary stochastic (R, S) policies

Recall that P_j^n represents the item-specific first period purchase decision, which is set to 1 if an order is issued in period 1, otherwise 0. Therefore, Constraints (4.50) guarantee the group fixed order cost in period 1 is properly counted. Constraints (4.51) ensure that the group fixed order cost is encountered whenever any item is replenished in period $2, \dots, T$. Constraints (4.52) ensure that there is no more than one outgoing arc from period 1. Constraints (4.53) are flow balance equations. Constraints (4.54) guarantee that period T is included in a replenishment cycle. By solving Fig. 4.7, the group order decision δ_t^n and item-specific order decision y_t^n are obtained,² for $t = 1, \dots, T$, $n = 1, \dots, N$.

4.B Computational results

This section presents detailed computational results of the (R, S) policy by adopting the shortest path reformulation discussed in Section 4.A, and compares that against (Q, S, T) policy (Özkaya et al., 2006), $Q(s, S)$ policy (Nielsen and Larsen, 2005), $P(s, S)$ policy (Viswanathan, 1997), (Q, S) policy (Pantumsinchai, 1992), MP policy (Atkins and Iyogun, 1988), $(s, c, S)_M$ policy (Melchiors, 2002), and $(s, c, S)_F$ policy (Federgruen et al., 1984), on data sets of Atkins and Iyogun

²This can be obtained by adding constraints $y_1^n = \sum_{j=1}^T Y_{1j}^n P_j^n$ and $y_i^n = \sum_{j=2}^T Y_{ij}^n$, $i = 2, \dots, T$, to Fig. 4.7.

(1988) and Viswanathan (1997) in Table 4.5 and 4.6. Note that “frequency” denotes the replenishment frequency.

K	b	h	(R, S)	average cost improvement $\Delta\%$							frequency
				(Q, S, T)	$Q(s, S)$	$P(s, S)$	(Q, S)	MP	$(s, c, S)_M$	$(s, c, S)_F$	
20	30	2	907.02	-3.10	-2.75	-2.30	5.18	-1.30	-3.87	-3.75	0.07
50	30	2	936.94	-0.91	-0.84	-0.33	4.38	0.68	0.79	2.14	0.06
100	30	2	990.50	-0.05	-0.45	0.75	2.57	1.77	4.39	6.81	0.05
150	30	2	1046.56	-0.24	-1.01	-0.35	0.52	0.65	5.68	8.36	0.05
200	30	2	1072.97	1.32	0.47	1.11	1.34	2.12	8.34	12.31	0.04
100	30	6	1639.75	-0.23	-1.52	-1.02	2.15	0.00	1.24	3.31	0.08
150	30	6	1707.05	0.64	-0.60	-0.07	1.46	0.95	2.34	6.68	0.07
200	30	6	1766.38	1.16	0.08	0.65	1.17	1.67	3.08	9.04	0.06
20	30	20	2388.07	-3.91	1.79	-2.63	5.82	-1.55	-3.11	-2.85	0.16
50	30	20	2469.10	-2.98	1.55	-1.69	4.49	-0.63	-1.40	-0.29	0.14
100	30	20	2596.50	-2.41	1.07	-1.34	2.80	-0.26	-0.99	3.49	0.13
150	30	20	2718.47	0.77	4.32	-1.26	1.27	-0.21	-0.59	6.20	0.10
200	30	20	2812.52	-3.23	0.14	-0.72	0.77	0.34	0.25	8.34	0.09

Table 4.5: Computational results on the data set of Atkins and Iyogun (1988)

K	b	h	(R, S)	average cost improvement $\Delta\%$							frequency
				(Q, S, T)	$Q(s, S)$	$P(s, S)$	(Q, S)	MP	$(s, c, S)_M$	$(s, c, S)_F$	
20	10	2	772.25	-0.03	0.48	0.76	8.30	1.79	-	1.80	0.06
50	10	2	813.94	-0.48	0.12	0.62	0.47	1.64	-	1.74	0.05
100	10	2	861.05	0.23	0.70	1.17	3.68	2.20	-	2.38	0.05
200	10	2	932.86	1.62	1.83	2.38	2.88	3.42	-	3.73	0.04
500	10	2	1131.42	0.14	0.14	0.59	0.18	1.60	-	2.12	0.02
20	10	6	1166.06	0.85	2.84	0.01	7.99	1.08	-	1.04	0.10
50	10	6	1222.82	-0.15	1.83	0.62	5.53	1.68	-	1.73	0.08
100	10	6	1283.92	1.33	2.50	1.26	4.49	2.34	-	2.46	0.07
200	10	6	1413.72	0.30	1.23	1.02	1.82	2.10	-	2.33	0.05
500	10	6	1658.48	2.26	2.20	2.52	2.30	3.59	-	4.03	0.04
20	10	10	1357.89	3.17	7.01	-0.05	8.02	1.06	-	0.98	0.11
50	10	10	1420.63	1.57	5.30	-0.03	5.88	1.07	-	1.07	0.09
100	10	10	1497.96	1.67	4.28	0.75	4.37	1.87	-	1.93	0.08
200	10	10	1637.27	0.66	2.18	1.15	2.16	2.28	-	2.44	0.06
500	10	10	1935.07	1.60	1.60	1.79	1.60	2.90	-	3.27	0.05
20	50	2	949.62	-4.70	-2.79	-2.16	5.46	-1.16	-	-1.18	0.08
50	50	2	1002.10	-4.80	-2.92	-2.32	2.28	-1.33	-	-1.27	0.06
100	50	2	1043.31	-1.95	-0.79	-0.23	1.98	0.78	-	0.92	0.05
200	50	2	1132.61	-1.29	-0.48	0.30	0.50	1.31	-	1.97	0.04
500	50	2	1327.95	0.08	0.08	0.82	0.13	1.83	-	2.30	0.03
20	50	6	1635.64	-2.97	-3.96	-3.53	4.58	-2.52	-	-2.55	0.12
50	50	6	1701.44	-1.91	-3.08	-2.61	2.84	-1.62	-	-1.57	0.09
100	50	6	1794.60	-1.37	-2.65	-2.09	0.94	-1.09	-	-0.97	0.09
200	50	6	1938.25	-0.27	-1.56	-0.89	-0.05	0.13	-	0.34	0.07
500	50	6	2244.01	-0.27	-0.27	0.43	-0.26	1.44	-	1.87	0.05
20	50	10	2056.14	-6.72	-3.40	-3.16	5.26	-2.14	-	-2.17	0.14
50	50	10	2131.28	-5.78	-2.60	-2.09	3.80	-1.08	-	-1.05	0.12
100	50	10	2237.94	-4.87	-1.91	-1.43	2.12	-0.41	-	-0.31	0.11
200	50	10	2448.79	-3.83	-2.11	-1.55	-0.75	-0.53	-	-0.34	0.08
500	50	10	2796.29	0.35	0.35	0.97	0.35	2.00	-	2.40	0.06
20	100	2	1015.96	-5.02	-4.31	-3.43	4.28	-2.43	-	-2.47	0.08
50	100	2	1054.36	-3.16	-2.49	-1.94	2.77	-0.95	-	-0.91	0.06
100	100	2	1114.82	-2.67	-2.03	-1.34	0.84	-0.34	-	-0.23	0.06
200	100	2	1200.38	-1.61	-0.94	-0.11	-0.01	0.90	-	1.13	0.05
500	100	2	1406.67	-0.76	-0.83	0.16	0.16	1.17	-	1.60	0.03
20	100	6	1806.52	-4.40	-5.30	-5.00	3.82	-4.01	-	-4.05	0.14
50	100	6	1863.01	-2.25	-3.27	-2.73	3.87	-1.74	-	-1.71	0.11
100	100	6	1965.83	-1.72	-2.90	-2.22	0.80	-1.22	-	-1.13	0.09
200	100	6	2106.78	0.44	-1.23	-0.48	0.94	0.54	-	0.73	0.07
500	100	6	2449.51	-0.88	-0.88	-0.07	-0.07	0.94	-	1.33	0.05
20	100	10	2326.30	-6.76	-4.75	-4.51	4.45	-3.51	-	-3.55	0.14
50	100	10	2414.57	-5.74	-3.67	-3.18	2.74	-2.18	-	-2.16	0.14
100	100	10	2527.82	-4.82	-2.80	-2.16	1.45	-1.15	-	-1.07	0.11
200	100	10	2728.08	-3.41	-1.90	-1.29	-0.49	-0.27	-	-0.10	0.09
500	100	10	3108.05	0.22	0.22	0.94	0.94	1.96	-	2.33	0.06
20	200	2	1074.45	-6.18	-5.53	-4.32	3.53	-3.33	-	-3.38	0.08

Continued on next page

Table 4.6 – continued from previous page

K	b	h	(R, S)	average cost improvement $\Delta\%$							frequency
				(Q, S, T)	Q(s, S)	P(s, S)	(Q, S)	MP	(s, c, S) _M	(s, c, S) _F	
50	200	2	1106.76	-3.50	-2.86	-2.16	2.85	-1.16	-	-1.14	0.06
100	200	2	1166.92	-2.56	-1.96	-1.28	0.99	-0.28	-	-0.18	0.06
200	200	2	1257.42	-1.62	-1.14	-0.19	-0.09	0.82	-	1.03	0.05
500	200	2	1470.29	-0.90	-0.90	0.05	0.05	1.05	-	1.45	0.03
20	200	6	1952.00	-5.02	-5.85	-5.53	3.73	-4.54	-	-4.59	0.13
50	200	6	2006.29	-2.56	-3.37	-2.70	3.04	-1.70	-	-1.69	0.11
100	200	6	2114.10	-1.66	-2.84	-2.07	0.95	-1.07	-	-0.99	0.09
200	200	6	2307.10	-1.65	-3.09	-2.26	-1.34	-1.26	-	-1.10	0.08
500	200	6	2620.77	-0.91	-0.91	0.08	0.08	1.09	-	1.45	0.05
20	200	10	2585.83	-7.26	-6.15	-5.83	3.39	-4.83	-	-4.88	0.15
50	200	10	2660.60	-5.21	-4.04	-3.48	2.51	-2.49	-	-2.48	0.15
100	200	10	2795.72	-4.68	-3.65	-2.87	0.61	-1.87	-	-1.81	0.11
200	200	10	2993.97	-3.17	-2.26	-1.37	-0.78	-0.36	-	-0.21	0.09
500	200	10	3421.28	-0.94	-0.94	-0.04	-0.04	0.97	-	1.30	0.07
20	1000	200	19742.15	-7.94	-5.39	-5.33	0.55	-4.32	-6.13	-4.39	0.50
50	1000	200	20093.33	-7.30	-5.08	-4.77	-0.00	-3.76	-	-3.81	0.50
100	1000	200	20557.35	-6.61	-4.72	-4.29	-0.75	-3.27	-5.03	-3.31	0.50
200	1000	200	21241.56	-5.85	-4.00	-3.43	-1.30	-2.42	-	-2.40	0.33
500	1000	200	22847.92	-4.07	-2.94	-2.32	-1.84	-1.31	-2.73	-1.22	0.25
20	1000	600	36291.82	-5.74	-3.34	-5.74	-1.12	-4.70	-5.45	-4.80	1.00
50	1000	600	36948.40	-6.36	-4.15	-5.68	-2.19	-4.66	-	-4.73	1.00
100	1000	600	37339.42	-5.90	-3.87	-4.67	-1.99	-3.63	-5.27	-3.69	0.67
200	1000	600	40175.00	-10.25	-8.40	-8.77	-7.13	-7.79	-	-7.81	0.66
500	1000	600	40150.15	-4.05	-2.86	-2.56	-2.16	-1.52	-2.85	-1.49	0.33
20	1000	1000	46309.40	-3.89	-2.82	-6.15	-2.40	-5.11	-6.62	-5.21	1.00
50	1000	1000	46886.02	-4.10	-3.36	-5.70	-2.97	-4.67	-	-4.76	1.00
100	1000	1000	47878.62	-5.13	-4.19	-6.07	-4.00	-5.03	-6.53	-5.11	1.00
200	1000	1000	48371.67	-4.34	-3.46	-4.46	-3.22	-3.42	-	-3.47	0.51
500	1000	1000	51354.21	-5.40	-4.44	-4.66	-4.26	-3.62	-4.93	-3.63	0.50
20	5000	200	28377.00	-10.13	-6.96	-6.73	-1.13	-5.74	-7.62	-5.80	1.00
50	5000	200	28692.13	-9.12	-6.00	-5.61	-1.37	-4.63	-	-4.67	0.60
100	5000	200	29200.60	-8.29	-5.45	-4.84	-1.79	-3.85	-6.06	-3.88	0.47
200	5000	200	29989.61	-7.39	-4.72	-4.09	-2.26	-3.10	-	-3.09	0.40
500	5000	200	31630.17	-4.75	-2.92	-2.08	-1.79	-1.09	-3.28	-1.02	0.25
20	5000	600	61029.16	-6.88	-5.71	-5.39	-1.71	-4.40	-5.99	-4.45	1.00
50	5000	600	61655.38	-6.69	-5.44	-5.16	-2.03	-4.17	-	-4.21	1.00
100	5000	600	62421.80	-6.14	-5.10	-4.73	-2.73	-3.74	-5.32	-3.77	0.60
200	5000	600	63682.80	-5.59	-4.69	-4.23	-2.98	-3.23	-	-3.25	0.59
500	5000	600	66527.80	-4.78	-4.05	-3.38	-3.29	-2.39	-4.14	-2.36	0.49
20	5000	1000	85618.15	-8.09	-5.53	-5.28	-2.62	-4.28	-5.87	-4.33	1.00
50	5000	1000	86331.93	-7.89	-5.23	-5.01	-2.92	-4.02	-	-4.06	1.00
100	5000	1000	87138.18	-6.90	-4.85	-4.52	-2.99	-3.52	-5.19	-3.56	1.00
200	5000	1000	89113.52	-6.98	-5.07	-4.56	-3.80	-3.57	-	-3.59	1.00
500	5000	1000	91330.15	-4.69	-3.44	-2.80	-2.51	-1.80	-3.58	-1.79	0.50
20	10000	200	31507.52	-9.31	-6.49	-6.25	-0.72	-5.26	-7.06	-5.32	0.67
50	10000	200	31897.92	-8.60	-5.78	-5.41	-1.25	-4.43	-	-4.46	0.65
100	10000	200	32101.59	-6.61	-4.23	-3.72	-0.54	-2.72	-4.82	-2.75	0.50
200	10000	200	33213.78	-6.61	-4.31	-3.62	-1.98	-2.63	-	-2.62	0.35
500	10000	200	35110.43	-4.90	-3.10	-2.27	-2.18	-1.28	-3.60	-1.22	0.25
20	10000	600	71079.04	-5.62	-5.33	-4.91	-1.77	-3.92	-5.59	-3.97	1.00
50	10000	600	71700.59	-5.29	-4.97	-4.63	-2.25	-3.64	-	-3.68	1.00
100	10000	600	72599.59	-5.01	-4.77	-4.35	-2.62	-3.35	-5.01	-3.38	1.00
200	10000	600	74654.19	-5.51	-5.34	-4.92	-3.69	-3.94	-	-3.95	1.00
500	10000	600	77718.34	-4.78	-4.71	-3.97	-3.96	-2.99	-4.84	-2.97	0.49
20	10000	1000	102578.36	-6.10	-5.05	-4.70	-2.42	-3.71	-5.38	-3.75	1.00
50	10000	1000	103239.08	-5.72	-4.67	-4.34	-2.52	-3.35	-	-3.38	1.00
100	10000	1000	104099.67	-5.21	-4.22	-3.89	-2.45	-2.89	-4.56	-2.92	1.00
200	10000	1000	106197.20	-5.25	-4.35	-3.88	-3.30	-2.88	-	-2.91	1.00
500	10000	1000	110044.84	-4.94	-4.18	-3.54	-3.44	-2.55	-4.44	-2.54	0.67
20	20000	200	34612.69	-8.16	-6.54	-6.31	-1.07	-5.33	-7.16	-5.38	1.00
50	20000	200	34853.07	-6.94	-5.45	-5.08	-1.00	-4.09	-	-4.13	0.60
100	20000	200	35734.53	-7.04	-5.70	-5.23	-2.39	-4.25	-6.53	-4.27	0.60
200	20000	200	36354.40	-5.59	-4.33	-3.61	-2.07	-2.62	-	-2.63	0.40
500	20000	200	38670.01	-4.80	-3.99	-3.12	-2.93	-2.14	-4.62	-2.09	0.38
20	20000	600	80553.66	-5.34	-5.07	-4.51	-1.65	-3.52	-5.38	-3.57	1.00
50	20000	600	81237.00	-5.05	-4.74	-4.24	-2.04	-3.26	-	-3.29	1.00
100	20000	600	82250.96	-4.75	-4.55	-4.05	-2.51	-3.06	-4.77	-3.08	1.00
200	20000	600	83710.17	-4.38	-4.19	-3.76	-2.79	-2.77	-	-2.78	0.67
500	20000	600	87509.30	-4.48	-4.31	-3.71	-3.61	-2.73	-4.25	-2.70	0.67
20	20000	1000	119658.22	-6.30	-5.36	-4.98	-2.98	-3.99	-5.61	-4.04	1.00
50	20000	1000	120589.66	-6.28	-5.18	-4.83	-3.21	-3.85	-	-3.88	1.00
100	20000	1000	121241.99	-5.63	-4.57	-4.16	-2.91	-3.18	-4.86	-3.20	1.00
200	20000	1000	123056.41	-5.16	-4.30	-3.90	-3.14	-2.92	-	-2.93	0.67

Continued on next page

Table 4.6 – continued from previous page

K	b	h	(R, S)	average cost improvement $\Delta\%$						frequency	
				(Q, S, T)	$Q(s, S)$	$P(s, S)$	(Q, S)	MP	$(s, c, S)_M$		$(s, c, S)_F$
500	20000	1000	127225.38	-4.71	-4.07	-3.54	-3.35	-2.55	-4.42	-2.55	0.67

Table 4.6: Computational results on the data set of Viswanathan (1997)

Bibliography

- Aggarwal, S. C. (1974). A review of current inventory theory and its applications. *International Journal of Production Research*, 12(4):443–482.
- Aneja, Y. and Noori, A. H. (1987). The optimality of (s, S) policies for a stochastic inventory problem with proportional and lump-sum penalty cost. *Management Science*, 33(6):750–755.
- Archibald, B. and Silver, E. (1978). (s, S) policies under continuous review and discrete compound poisson demand. *Management Science*, 24(9):899–909.
- Arkin, E., Joneja, D., and Roundy, R. (1989). Computational complexity of uncapacitated multi-echelon production planning problems. *Operations research letters*, 8(2):61–66.
- Arrow, K. J., Harris, T., and Marschak, J. (1951). Optimal inventory policy. *Econometrica: Journal of the Econometric Society*, pages 250–272.
- Askin, R. G. (1981). A procedure for production lot sizing with probabilistic dynamic demand. *AIIE Transactions*, 13(2):132–137.
- Atkins, D. R. and Iyogun, P. O. (1988). Periodic versus “can-order” policies for coordinated multi-item inventory systems. *Management Science*, 34(6):791–796.
- Axsäter, S. (2010). *Inventory Control (International Series in Operations Research & Management Science)*. Springer, 2nd ed. edition.
- Baker, K. R. (1989). Lot-sizing procedures and a standard data set: a reconciliation of the literature. *Journal of Manufacturing and Operations Management*, 2(3):199–221.
- Balintfy, J. L. (1964). On a basic class of multi-item inventory problems. *Management Science*, 10(2):287–297.

- Bashyam, S. and Fu, M. C. (1998). Optimization of (s, S) inventory systems with random lead times and a service level constraint. *Management Science*, 44(12-part-2):S243–S256.
- Bastos, L. D. S. L., Mendes, M. L., Nunes, D. R. D. L., Melo, A. C. S., and Carneiro, M. P. (2017). A systematic literature review on the joint replenishment problem solutions: 2006-2015. *Production*, 27.
- Bellman, R. (1966). Dynamic programming. *Science*, 153(3731):34–37.
- Bellman, R. E. and Dreyfus, S. E. (2015). *Applied dynamic programming*, volume 2050. Princeton university press.
- Belotti, P., Bonami, P., Fischetti, M., Lodi, A., Monaci, M., Nogales-Gómez, A., and Salvagnin, D. (2016). On handling indicator constraints in mixed integer programming. *Computational Optimization and Applications*, 65(3):545–566.
- Bertsimas, D. and Thiele, A. (2006). A robust optimization approach to inventory theory. *Operations research*, 54(1):150–168.
- Beyer, D. and Sethi, S. P. (1997). Average cost optimality in inventory models with markovian demands. *Journal of Optimization Theory and Applications*, 92(3):497–526.
- Beyer, D., Sethi, S. P., and Taksar, M. (1998). Inventory models with markovian demands and cost functions of polynomial growth. *Journal of Optimization theory and Applications*, 98(2):281–323.
- Beyer, H. G. and Sendhoff, B. (2007). Robust optimization—a comprehensive survey. *Computer methods in applied mechanics and engineering*, 196(33-34):3190–3218.
- Billingsley, P. (2008). *Probability and measure*. John Wiley & Sons.
- Birge, J. R. and Louveaux, F. (2011). *Introduction to stochastic programming*. Springer Science & Business Media.
- Blackburn, J. D. and Millen, R. A. (1980). Heuristic lot-sizing performance in a rolling-schedule environment. *Decision Sciences*, 11(4):691–701.
- Boctor, F. F., Laporte, G., and Renaud, J. (2004). Models and algorithms for the dynamic-demand joint replenishment problem. *International Journal of Production Research*, 42(13):2667–2678.

- Bollapragada, S. and Morton, T. E. (1999). A simple heuristic for computing nonstationary (s, S) policies. *Operations Research*, 47(4):576–584.
- Bookbinder, J. H. and Tan, J. Y. (1988). Strategies for the probabilistic lot-sizing problem with service-level constraints. *Management Science*, 34(9):1096–1108.
- Box, G. E., Jenkins, G. M., Reinsel, G. C., and Ljung, G. M. (2015). *Time series analysis: forecasting and control*. John Wiley & Sons.
- Brahimi, N., Dauzere-Peres, S., Najid, N. M., and Nordli, A. (2006). Single item lot sizing problems. *European Journal of Operational Research*, 168(1):1–16.
- Brown, R. G. (1967). *Decision rules for inventory management*. New York: Holt, Rinehart and Winston.
- Bushuev, M. A., Guiffrida, A., Jaber, M., and Khan, M. (2015). A review of inventory lot sizing review papers. *Management Research Review*, 38(3):283–298.
- Carrizosa, E., Olivares-Nadal, A. V., and Ramírez-Cobo, P. (2016). Robust newsvendor problem with autoregressive demand. *Computers & Operations Research*, 68:123–133.
- Chen, F. and Song, J. S. (2001). Optimal policies for multiechelon inventory problems with markov-modulated demand. *Operations Research*, 49(2):226–234.
- Cheng, F. and 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(2):183–192.
- Cobb, B. R. (2016). Inventory control for returnable transport items in a closed-loop supply chain. *Transportation Research Part E: Logistics and Transportation Review*, 86:53–68.
- Cohen, M. A., Kleindorfer, P. R., and Lee, H. L. (1988). Service constrained (s, S) inventory systems with priority demand classes and lost sales. *Management Science*, 34(4):482–499.
- Datta, T. and Pal, A. (1991). Effects of inflation and time-value of money on an inventory model with linear time-dependent demand rate and shortages. *European Journal of Operational Research*, 52(3):326–333.

- De Kok, T., Inderfurth, K., et al. (1997). Nervousness in inventory management: comparison of basic control rules. *European Journal of Operational Research*, 103(1):55–82.
- Dobson, G., Pinker, E. J., and Yildiz, O. (2017). An EOQ model for perishable goods with age-dependent demand rate. *European Journal of Operational Research*, 257(1):84–88.
- Dong, L. and Lee, H. L. (2003). Optimal policies and approximations for a serial multiechelon inventory system with time-correlated demand. *Operations Research*, 51(6):969–980.
- Drake, M. J. and Marley, K. A. (2014). A century of the EOQ. In *Handbook of EOQ Inventory Problems*, pages 3–22. Springer.
- Dural-Selcuk, G., Kilic, O., Tarim, S., and Rossi, R. (2016). A comparison of non-stationary stochastic lot-sizing strategies. *arXiv:1607.08896*.
- Dvoretzky, A., Kiefer, J., and Wolfowitz, J. (1952). The inventory problem: I. case of known distributions of demand. *Econometrica (pre-1986)*, 20(2):187.
- Ehrhardt, R. (1979). The power approximation for computing (s, S) inventory policies. *Management Science*, 25(8):777–786.
- Elms, D. K. and Low, P. (2013). *Global value chains in a changing world*. World Trade Organization Geneva.
- Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica: Journal of the Econometric Society*, pages 987–1007.
- Fabens, A. and Karlin, S. (1960). A Stationary Inventory Model with Markovian Demand. In *Mathematical Methods in the Social Sciences, 1959: Proceedings*, number 4, page 159. Stanford University Press.
- Federgruen, A., Groenevelt, H., and Tijms, H. C. (1984). Coordinated replenishments in a multi-item inventory system with compound poisson demands. *Management Science*, 30(3):344–357.
- Federgruen, A. and Zipkin, P. (1984). An efficient algorithm for computing optimal (s, S) policies. *Operations research*, 32(6):1268–1285.

- Feng, Y. and Xiao, B. (2000). A new algorithm for computing optimal (s, S) policies in a stochastic single item/location inventory system. *IIE Transactions*, 32(11):1081–1090.
- Fotopoulos, S., Wang, M.-C., and Rao, S. S. (1988). Safety stock determination with correlated demands and arbitrary lead times. *European Journal of Operational Research*, 35(2):172–181.
- Frauendorfer, K. (1996). Barycentric scenario trees in convex multistage stochastic programming. *Mathematical Programming*, 75(2):277–293.
- Fung, R. and Ma, X. (2001). A new method for joint replenishment problems. *Journal of the Operational Research Society*, 52(3):358–362.
- Gallego, G. and Sethi, S. (2005). K-Convexity in \mathcal{R}^n . *Journal of Optimization Theory and Applications*, 127(1):71–88.
- Glock, C. H., Grosse, E. H., and Ries, J. M. (2014). The lot sizing problem: A tertiary study. *International Journal of Production Economics*, 155:39–51.
- Goyal, S. K. and Belton, A. S. (1979). Note ON “A SIMPLE METHOD OF DETERMINING ORDER QUANTITIES IN JOINT REPLENISHMENTS UNDER DETERMINISTIC DEMAND”. *Management Science (pre-1986)*, 25(6):604.
- Goyal, S. K. and Deshmukh, S. G. (1993). Discussion A note on ‘The economic ordering quantity for jointly replenishing items’. *International Journal of Production Research*, 31(12):2959–2961.
- Goyal, S. K. and Satir, A. T. (1989). Joint replenishment inventory control: deterministic and stochastic models. *European Journal of Operational Research*, 38(1):2–13.
- Graves, S. (1999). A single-item inventory model for a nonstationary demand process. *Manufacturing & Service Operations Management*, 1(1):50–61.
- Gutierrez-Alcoba, A., Rossi, R., Martin-Barragan, B., and Hendrix, E. M. (2016). A simple heuristic for perishable item inventory control under non-stationary stochastic demand. *International Journal of Production Research*, 55(7):1885–1897.
- Hariga, M. (1994). Two new heuristic procedures for the joint replenishment problem. *Journal of the Operational Research Society*, 45(4):463–471.

- Harris, F. W. (1913). How many parts to make at once. *Factory, The Magazine of Management*, 10(2):135–136.
- Heisig, G. (1998). Planning stability under (s, S) inventory control rules. *Operations-Research-Spektrum*, 20(4):215–228.
- Heisig, G. (2001). Comparison of (s, S) and (s, nQ) inventory control rules with respect to planning stability. *International Journal of Production Economics*, 73(1):59–82.
- Hou, K.-L. (2006). An inventory model for deteriorating items with stock-dependent consumption rate and shortages under inflation and time discounting. *European Journal of Operational Research*, 168(2):463–474.
- Hu, J., Zhang, C., and Zhu, C. (2016). (s, S) inventory systems with correlated demands. *INFORMS Journal on Computing*, 28(4):603–611.
- Hua, Z., Yang, J., Huang, F., and Xu, X. (2009). A static-dynamic strategy for spare part inventory systems with nonstationary stochastic demand. *Journal of the Operational Research Society*, 60(9):1254–1263.
- IBM (2011). *IBM ILOG CPLEX Optimization Studio OPL Language Reference Manual*.
- Iglehart, D. (1962). *Optimal Policy for Dynamic Inventory Process With Nonstationary Stochastic Demands. Chapter 8 in Studies in Applied Probability and Management Science, K. Arrow, S. Karlin, & H. Scarf*. Stanford University Press, Stanford, Calif.
- Iglehart, D. L. (1963). Optimality of (s, S) policies in the infinite horizon dynamic inventory problem. *Management science*, 9(2):259–267.
- Ignall, E. (1969). Optimal continuous review policies for two product inventory systems with joint setup costs. *Management Science*, 15(5):278–283.
- Janssen, F. and de Kok, T. (1999). A two-supplier inventory model. *International journal of production economics*, 59(1-3):395–403.
- Johansen, S. G. and Melchior, P. (2003). Can-order policy for the periodic-review joint replenishment problem. *Journal of the Operational Research Society*, 54(3):283–290.

- Johnson, E. L. (1967). Optimality and computation of (σ, S) policies in the multi-item infinite horizon inventory problem. *Management Science*, 13(7):475–491.
- Johnson, G. D. and Thompson, H. (1975). Optimality of myopic inventory policies for certain dependent demand processes. *Management Science*, 21(11):1303–1307.
- Jünger, M., Liebling, T. M., Naddef, D., Nemhauser, G. L., Pulleyblank, W. R., Reinelt, G., Rinaldi, G., and Wolsey, L. A. (2009). *50 Years of integer programming 1958-2008: From the early years to the state-of-the-art*. Springer Science & Business Media.
- Kalin, D. (1980). On the optimality of (σ, S) policies. *Mathematics of Operations Research*, 5(2):293–307.
- Kall, P., Wallace, S. W., and Kall, P. (1994). *Stochastic programming*. Springer.
- Kalpakam, S. and Arivarignan, G. (1988). A continuous review perishable inventory model. *Statistics*, 19(3):389–398.
- Kalpakam, S. and Arivarignan, G. (1993). A coordinated multicommodity (s, S) inventory system. *Mathematical and Computer Modelling*, 18(11):69–73.
- Kalpakam, S. and Sapna, K. (1994). Continuous review (s, S) inventory system with random lifetimes and positive leadtimes. *Operations Research Letters*, 16(2):115–119.
- Kaspi, M. and Rosenblatt, M. J. (1983). An improvement of silver’s algorithm for the joint replenishment problem. *AIIE Transactions*, 15(3):264–267.
- Kaspi, M. and Rosenblatt, M. J. (1991). On the economic ordering quantity for jointly replenished items. *International Journal of Production Research*, 29(1):107–114.
- Kayış, E., Bilgiç, T., and Karabulut, D. (2008). A note on the can-order policy for the two-item stochastic joint-replenishment problem. *IIE Transactions*, 40(1):84–92.
- Khouja, M. and Goyal, S. (2008). A review of the joint replenishment problem literature: 1989–2005. *European Journal of Operational Research*, 186(1):1–16.
- Kilic, O. A. and Tarim, S. A. (2011). An investigation of setup instability in non-stationary stochastic inventory systems. *International Journal of Production Economics*, 133(1):286–292.

- Kuhn, D. (2006). *Generalized bounds for convex multistage stochastic programs*, volume 548. Springer Science & Business Media.
- Kwon, W. H. and Han, S. H. (2006). *Receding horizon control: model predictive control for state models*. Springer Science & Business Media.
- Leenaerts, D. and Van Bokhoven, W. M. (2013). *Piecewise linear modeling and analysis*. Springer Science & Business Media.
- Li, X. (2013). Managing dynamic inventory systems with product returns: a markov decision process. *Journal of Optimization Theory and Applications*, 157(2):577–592.
- Lian, Z., Liu, X., and Zhao, N. (2009). A perishable inventory model with markovian renewal demands. *International Journal of Production Economics*, 121(1):176–182.
- Liberatore, M. J. (1979). The eoq model under stochastic lead time. *Operations Research*, 27(2):391–396.
- Lim, Y. F. and Wang, C. (2017). Inventory management based on target-oriented robust optimization. *Management Science*, 63(12):4409–4427.
- Liu, B. and Esogbue, A. O. (2012). *Decision criteria and optimal inventory processes*, volume 20. Springer Science & Business Media.
- Liu, L. (1990). (s, S) continuous review models for inventory with random lifetimes. *Operations Research Letters*, 9(3):161–167.
- Liu, L. and Lian, Z. (1999). (s, S) continuous review models for products with fixed lifetimes. *Operations Research*, 47(1):150–158.
- Liu, L. and Shi, D. H. (1999). An (s, S) model for inventory with exponential lifetimes and renewal demands. *Naval Research Logistics (NRL)*, 46(1):39–56.
- Liu, L. and Yang, T. (1999). An (s, S) random lifetime inventory model with a positive lead time. *European Journal of Operational Research*, 113(1):52–63.
- Melchioris, P. (2002). Calculating can-order policies for the joint replenishment problem by the compensation approach. *European Journal of Operational Research*, 141(3):587–595.

- Nasr, W. W. and Maddah, B. (2015). Continuous (s, S) policy with mppp correlated demand. *European Journal of Operational Research*, 246(3):874–885.
- Nasri, F., affisco, J. F., and Javad Paknejad, M. (1990). Setup cost reduction in an inventory model with finite-range stochastic lead times. *The International Journal of Production Research*, 28(1):199–212.
- Natarajan, K. and Teo, C. P. (2017). On reduced semidefinite programs for second order moment bounds with applications. *Mathematical Programming*, 161(1-2):487–518.
- Nielsen, C. and Larsen, C. (2005). An analytical study of the $Q(s, S)$ policy applied to the joint replenishment problem. *European Journal of Operational Research*, 163(3):721–732.
- Nilsson, A., Segerstedt, A., and Van Der Sluis, E. (2007). A new iterative heuristic to solve the joint replenishment problem using a spreadsheet technique. *International Journal of Production Economics*, 108(1-2):399–405.
- Ohno, K. and Ishigaki, T. (2001). A multi-item continuous review inventory system with compound poisson demands. *Mathematical Methods of Operations Research*, 53(1):147–165.
- Ohno, K., Ishigaki, T., and Yoshii, T. (1994). A new algorithm for a multi-item periodic review inventory system. *Zeitschrift für Operations Research*, 39(3):349–364.
- Özekici, S. and Parlar, M. (1999). Inventory models with unreliable suppliers in a random environment. *Annals of Operations Research*, 91:123–136.
- Özen, U., Doğru, M. K., and Tarim, S. A. (2012). Static-dynamic uncertainty strategy for a single-item stochastic inventory control problem. *Omega*, 40(3):348–357.
- Özkaya, B. Y., Gürler, Ü., and Berk, E. (2006). The stochastic joint replenishment problem: A new policy, analysis, and insights. *Naval Research Logistics (NRL)*, 53(6):525–546.
- Padmanabhan, G. and Vrat, P. (1995). EOQ models for perishable items under stock dependent selling rate. *European Journal of Operational Research*, 86(2):281–292.

- Pantumsinchai, P. (1992). A comparison of three joint ordering inventory policies. *Decision Sciences*, 23(1):111–127.
- Parlar, M. and Berkin, D. (1991). Future supply uncertainty in eoq models. *Naval Research Logistics (NRL)*, 38(1):107–121.
- Pauls-Worm, K. G., Hendrix, E. M., Haijema, R., and van der Vorst, J. G. (2014). An MILP approximation for ordering perishable products with non-stationary demand and service level constraints. *International Journal of Production Economics*, 157:133–146.
- Pentico, D. W. and Drake, M. J. (2011). A survey of deterministic models for the EOQ and EPQ with partial backordering. *European Journal of Operational Research*, 214(2):179–198.
- Pervin, M., Roy, S. K., and Weber, G. W. (2018). Analysis of inventory control model with shortage under time-dependent demand and time-varying holding cost including stochastic deterioration. *Annals of Operations Research*, 260(1-2):437–460.
- Powell, W. B. (2007). *Approximate Dynamic Programming: Solving the curses of dimensionality*, volume 703. John Wiley & Sons.
- Powell, W. B. (2009). What you should know about approximate dynamic programming. *Naval Research Logistics (NRL)*, 56(3):239–249.
- Qiu, R., Sun, M., and Lim, Y. F. (2017). Optimizing (s, S) policies for multi-period inventory models with demand distribution uncertainty: Robust dynamic programming approaches. *European Journal of Operational Research*, 261(3):880–892.
- Ravichandran, N. (1995). Stochastic analysis of a continuous review perishable inventory system with positive lead time and poisson demand. *European Journal of Operational Research*, 84(2):444–457.
- Ray, W. (1981). Computation of reorder levels when the demands are correlated and the lead time random. *Journal of the Operational Research Society*, 32(1):27–34.
- Relvas, S., Magatão, S. N. B., Barbosa-Póvoa, A. P. F., and Neves Jr, F. (2013). Integrated scheduling and inventory management of an oil products distribution system. *Omega*, 41(6):955–968.

- Renberg, B. and Planche, R. (1967). Un modele pour la gestion simultanee des n articles d'un stock. *Revue Francaise d'Informatique et de Recherche Operationelle*, 6:47–59.
- Rossi, R., Kilic, O. A., and Tarim, S. A. (2015). Piecewise linear approximations for the static–dynamic uncertainty strategy in stochastic lot-sizing. *Omega*, 50:126–140.
- Rossi, R., Tarim, S. A., and Bollapragada, R. (2012a). Constraint-based local search for inventory control under stochastic demand and lead time. *INFORMS journal on computing*, 24(1):66–80.
- Rossi, R., Tarim, S. A., Hnich, B., and Prestwich, S. (2008). A global chance-constraint for stochastic inventory systems under service level constraints. *Constraints*, 13(4):490–517.
- Rossi, R., Tarim, S. A., Hnich, B., and Prestwich, S. (2010). Computing the non-stationary replenishment cycle inventory policy under stochastic supplier lead-times. *International Journal of Production Economics*, 127(1):180–189.
- Rossi, R., Tarim, S. A., Hnich, B., and Prestwich, S. (2012b). Constraint programming for stochastic inventory systems under shortage cost. *Annals of Operations Research*, 195(1):49–71.
- Rossi, R., Tarim, S. A., Prestwich, S., and Hnich, B. (2014). Piecewise linear lower and upper bounds for the standard normal first order loss function. *Applied Mathematics and Computation*, 231:489–502.
- Sahin, I. (1982). On the objective function behavior in (s, S) inventory models. *Operations Research*, 30(4):709–724.
- Salameh, M. K., Yassine, A. A., Maddah, B., and Ghaddar, L. (2014). Joint replenishment model with substitution. *Applied Mathematical Modelling*, 38(14):3662–3671.
- Scarf, H. (1960). Optimality of (s, S) policies in the dynamic inventory problem. In *Mathematical Methods in the Social Science*, pages 196–202. Stanford University Press.
- Schaack, J. P. and Silver, E. (1972). A procedure, involving simulation, for selecting the control variables of an (S, c, s) joint ordering strategy. *INFOR: Information Systems and Operational Research*, 10(2):154–170.

- Schneider, H. and Ringuest, J. L. (1990). Power approximation for computing (s, S) policies using service level. *Management Science*, 36(7):822–834.
- Sethi, S. P. and Cheng, F. (1997). Optimality of (s, S) policies in inventory models with markovian demand. *Operations Research*, 45(6):931–939.
- Si, J., Barto, A. G., Powell, W. B., and Wunsch, D. (2004). *Handbook of learning and approximate dynamic programming*, volume 2. John Wiley & Sons.
- Silver, E. A. (1974). A control system for coordinated inventory replenishment. *International Journal of Production Research*, 12(6):647–671.
- Silver, E. A. (1976). A simple method of determining order quantities in joint replenishments under deterministic demand. *Management Science*, 22(12):1351–1361.
- Silver, E. A. (1979). A simple inventory replenishment decision rule for a linear trend in demand. *Journal of the Operational Research society*, 30(1):71–75.
- Silver, E. A. (1981). Establishing reorder points in the (S, c, s) coordinated control system under compound poisson demand. *International Journal of Production Research*, 19(6):743–750.
- Silver, E. A. and Meal, H. C. (1973). A heuristic for selecting lot size quantities for the case of a deterministic time-varying demand rate and discrete opportunities for replenishment. *Production and Inventory Management Journal*, 2:64–74.
- Silver, E. A. and Peterson, R. (1985). *Decision systems for inventory management and production planning*. John Wiley & Sons Inc.
- Silver, E. A., Pyke, D. F., and Peterson, R. (1998). *Inventory management and production planning and scheduling*, volume 3. Wiley New York.
- Snyder, L. V. and Shen, Z. J. M. (2011). *Fundamentals of supply chain theory*. John Wiley & Sons.
- Song, J. S. and Zipkin, P. (1993). Inventory control in a fluctuating demand environment. *Operations Research*, 41(2):351–370.
- Stidham, S. (1977). Cost models for stochastic clearing systems. *Operations Research*, 25(1):100–127.

- Strijbosch, L., Syntetos, A., Boylan, J., and Janssen, E. (2011). On the interaction between forecasting and stock control: the case of non-stationary demand. *International Journal of Production Economics*, 133(1):470–480.
- Taboga, M. (2012). *Lectures on probability theory and mathematical statistics*. CreateSpace Independent Publishing Platform.
- Tarim, S. A., Dogru, M. K., Özen, U., and Rossi, R. (2011). An efficient computational method for a stochastic dynamic lot-sizing problem under service-level constraints. *European Journal of Operational Research*, 215(3):563–571.
- Tarim, S. A. and Kingsman, B. G. (2004). The stochastic dynamic production/inventory lot-sizing problem with service-level constraints. *International Journal of Production Economics*, 88(1):105–119.
- Tarim, S. A. and Kingsman, B. G. (2006). 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.
- Tarim, S. A. and Smith, B. M. (2008). Constraint programming for computing non-stationary (R, S) inventory policies. *European Journal of Operational Research*, 189(3):1004–1021.
- Tempelmeier, H. (2007). On the stochastic uncapacitated dynamic single-item lotsizing problem with service level constraints. *European Journal of Operational Research*, 181(1):184–194.
- Tempelmeier, H. (2013). Stochastic lot sizing problems. In *Handbook of Stochastic Models and Analysis of Manufacturing System Operations*, pages 313–344. Springer.
- Thompson, R. M. and Silver, E. A. (1975). A coordinated inventory control system for compound poisson demand and zero lead time. *International Journal of Production Research*, 13(6):581–602.
- Tunc, H., Kilic, O. A., Tarim, S. A., and Eksioglu, B. (2011). The cost of using stationary inventory policies when demand is non-stationary. *Omega*, 39(4):410–415.
- Tunc, H., Kilic, O. A., Tarim, S. A., and Eksioglu, B. (2013). A simple approach for assessing the cost of system nervousness. *International Journal of Production Economics*, 141(2):619–625.

- Tunc, H., Kilic, O. A., Tarim, S. A., and Eksioglu, B. (2014). A reformulation for the stochastic lot sizing problem with service-level constraints. *Operations Research Letters*, 42(2):161–165.
- Tunc, H., Kilic, O. A., Tarim, S. A., and Rossi, R. (2018). An extended mixed-integer programming formulation and dynamic cut generation approach for the stochastic lot-sizing problem. *INFORMS Journal on Computing*, 30(3):492–506.
- Ullah, H. and Parveen, S. (2010). A literature review on inventory lot sizing problems. *Global Journal of Research In Engineering*, 10(5):21–36.
- Van Eijs, M., Heuts, R. M. J., and Kleijnen, J. P. C. (1992). Analysis and comparison of two strategies for multi-item inventory systems with joint replenishment costs. *European Journal of Operational Research*, 59(3):405–412.
- Veinott Jr, A. F. (1965). Optimal policy for a multi-product, dynamic, nonstationary inventory problem. *Management Science*, 12(3):206–222.
- Veinott Jr, A. F. (1966). On the optimality of (s, S) inventory policies: New conditions and a new proof. *SIAM Journal on Applied Mathematics*, 14(5):1067–1083.
- Veinott Jr, A. F. and Wagner, H. M. (1965). Computing optimal (s, S) inventory policies. *Management Science*, 11(5):525–552.
- Viswanathan, S. (1996). A new optimal algorithm for the joint replenishment problem. *Journal of the Operational Research Society*, 47(7):936–944.
- Viswanathan, S. (1997). Note. Periodic review (s, S) policies for joint replenishment inventory systems. *Management Science*, 43(10):1447–1454.
- Viswanathan, S. (2002). On optimal algorithms for the joint replenishment problem. *Journal of the Operational Research Society*, 53(11):1286–1290.
- Wagner, H. M. and Whitin, T. M. (1958). Dynamic version of the economic lot size model. *Management science*, 5(1):89–96.
- Weiss, H. J. (1980). Optimal ordering policies for continuous review perishable inventory models. *Operations Research*, 28(2):365–374.
- Weiss, N. A. (2006). *A course in probability*. Addison-Wesley.

- Wildeman, R., Frenk, J., and Dekker, R. (1997). An efficient optimal solution method for the joint replenishment problem. *European Journal of Operational Research*, 99(2):433–444.
- Wolsey, L. A. (1998). *Integer programming*. Wiley.
- Wu, K. S., Ouyang, L. Y., and Yang, C. T. (2006). An optimal replenishment policy for non-instantaneous deteriorating items with stock-dependent demand and partial backlogging. *International Journal of Production Economics*, 101(2):369–384.
- Xu, H. and Wang, H. P. B. (1990). An economic ordering policy model for deteriorating items with time proportional demand. *European Journal of Operational Research*, 46(1):21–27.
- Xu, Y., Bisi, A., and Dada, M. (2010). New structural properties of (s, S) policies for inventory models with lost sales. *Operations Research Letters*, 38(5):441–449.
- Yano, C. A. and Lee, H. L. (1995). Lot sizing with random yields: A review. *Operations Research*, 43(2):311–334.
- Zheng, Y. S. and Federgruen, A. (1991). Finding optimal (s, S) policies is about as simple as evaluating a single policy. *Operations research*, 39(4):654–665.
- Zipkin, P. H. (2000). *Foundations of inventory management*. McGraw-Hill.