

Inventory control
for a perishable product
with non-stationary demand

K.G.J. Pauls-Worm

Thesis committee

Promotor

Prof. Dr J.G.A.J. van der Vorst
Professor of Logistics and Operations Research
Wageningen University

Co-promotors

Dr E.M.T. Hendrix
Profesor Contratado Doctor, Departamento de Arquitectura de Computadores,
Universidad de Málaga, Spain
Associate professor, Operations Research and Logistics Group
Wageningen University

Dr R. Haijema
Associate professor, Operations Research and Logistics Group
Wageningen University

Other members

Prof. Dr G.J.J.A.N. van Houtum, Eindhoven University of Technology
Prof. Dr S. Minner, Technische Universität München, Germany
Prof. Dr K.J. Roodbergen, University of Groningen
Prof. Dr A.G.J.M. Oude Lansink, Wageningen University

This research was conducted under the auspices of Wageningen School of Social Sciences (WASS)

Inventory control for a perishable product with non-stationary demand

K.G.J. Pauls-Worm

Thesis

submitted in fulfillment of the requirements for the degree of doctor
at Wageningen University
by the authority of the Rector Magnificus
Prof. Dr A.P.J. Mol,
in the presence of the
Thesis Committee appointed by the Academic Board
to be defended in public
on Tuesday 11th of October 2016
at 11 a.m. in the Aula.

K.G.J. Pauls-Worm

Inventory control for a perishable product with non-stationary demand
142 pages.

PhD thesis, Wageningen University, Wageningen, NL (2016)
With references, with summary in English

ISBN: 978-94-6257-850-0

DOI: 10.18174/385217

Contents

1	Introduction.....	1
1.1	Background	2
1.2	Inventory control	5
1.2.1	Key questions.....	5
1.2.2	Relevant costs.....	7
1.2.3	Service levels	8
1.2.4	Perishability in inventory literature.....	9
1.3	Problem statement.....	10
1.4	Research opportunities.....	10
1.5	Thesis outline	13
1.6	Overview of papers in this thesis	13
2	MILP approximation generating a YS policy	15
2.1	Introduction	17
2.2	Literature review.....	19
2.3	SP model for a perishable product	20
2.4	Deterministic MILP approximation.....	22
2.4.1	The waste-compensating replenishment cycle policy.....	22
2.4.2	Objective function	23
2.4.3	Service-level constraint	23
2.4.4	FIFO constraints.....	25
2.4.5	MILP model for a perishable product.....	26
2.5	Numerical illustration of the MILP model	27
2.5.1	Numerical illustration of the MILP model with FIFO issuance	28
2.5.2	FIFO relaxation	29
2.6	Results of the MILP model.....	30
2.6.1	Design of experiments and performance measures	31
2.6.2	Performance of the MILP approximation	33
2.6.3	Benchmark with SDP	37
2.7	Conclusions	38
	Appendix 2.A.....	40
	Appendix 2.B.....	41

3	SDP in inventory control: a $Q(X)$ policy.....	45
3.1	Introduction	47
3.2	Dynamic Programming	48
3.3	Stochastic Dynamic Programming	50
3.4	SDP with an α -service level constraint.....	51
3.5	Fill rate constraint	55
3.6	Continuous demand	57
3.7	Conclusions	60
4	Comparing order policies and solution methods	61
4.1	Introduction	63
4.2	Literature	64
4.3	Stochastic Programming model	65
4.4	Characteristics of a solution to the static-dynamic uncertainty strategy	67
4.4.1	Replenishment cycles and limits on timing vector Y	67
4.4.2	Basic order-up-to level and optimal order quantities	68
4.4.3	Monte Carlo estimation of the service level.....	69
4.5	YS policy: sample based SMC-MINLP approach	70
4.5.1	MC-MILP optimisation of the YS policy	70
4.5.2	Smoothed Monte Carlo MINLP approach to the YS policy	71
4.6	$YQ(X)$ policy: sample based algorithm	73
4.7	Flexible $Q(X)$ policy	75
4.8	Comparison of policies	76
4.9	Conclusion.....	78
5	YS policy in retail	81
5.1	Introduction	83
5.2	Model assumptions	85
5.3	Stochastic Programming model	85
5.4	Order policies	87
5.5	MILP model	90
5.6	Numerical experiments.....	92
5.6.1	Design of experiments	92
5.6.2	Results of 27 scenarios	92

5.6.3 Scenario in detail	94
5.7 Conclusion.....	95
6 MILP approximation generating a YQ policy	97
6.1 Introduction	99
6.2 Literature	100
6.3 Stochastic Programming Model	102
6.4 Towards a production plan	105
6.5 MILP model	106
6.5.1 MILP model formulation.....	106
6.5.2 Numerical illustration of the MILP model	108
6.6 Results of the MILP model.....	109
6.6.1 Design of Experiments	109
6.6.2 Benchmark with a scenario-based approach.....	110
6.6.3 Experiments	111
6.7 Conclusions	114
7 Conclusion and general discussion	115
7.1 Introduction.....	116
7.2 Conclusion.....	116
7.2.1 RO1: MILP approximation generating a YS policy	116
7.2.2 RO2: SDP in inventory control: a Q(X) policy	117
7.2.3 RO3: Comparing policies and solution methods	118
7.2.4 RO4: YS policy in retail	118
7.2.5 RO5: MILP approximation for a YQ policy.....	119
7.3 Integrated findings	119
7.3.1 Scientific contribution	120
7.3.2 Managerial impact	123
7.4 Discussion	125
7.5 Directions for future research	128
References.....	131
Summary	136
Acknowledgements.....	140

Chapter 1

Introduction

1.1 Background

Globally, around one-third of the edible parts of perishable food products is being thrown away by farms, factories, supermarkets, shops, restaurants and households, every year an estimated 1.3 billion ($1.3 \cdot 10^9$) ton (FAO, 2011). According to the Food and Agriculture Organization of the United Nations (FAO, 2015), *food loss is defined as "the decrease in quantity or quality of food"* and comprises *agricultural or fisheries products intended for human consumption that are ultimately not eaten by people or that have incurred a reduction in quality reflected in their nutritional value, economic value or food safety*. Food loss arises at production of raw material at the farm, at post-harvest handling and storage and at processing of the raw material (Parfitt et al., 2010). Food waste refers to food loss at the end of the food supply chain of food that was fit for human consumption and is mainly a result of retailer and consumer behaviour (Parfitt et al., 2010). Food waste occurs in two ways, either in markdowns when products are still saleable but approaching the end of their lifetime or appearing less attractive, or in garbage when products are no longer (re)saleable, useable or edible. In Europe the total food loss and waste is 31% of the initial production from which 6.1% occurs in the food processing, packaging and distribution (HLPE, 2014). Reducing the annual food loss and waste will result in benefits for companies, consumers and the environment in terms of money, volume, energy and sustainability. In this thesis the term *waste* refers to both food loss and food waste.

A way to reduce waste in the food supply chain is to control the inventory levels in the supply chain. Inventory control has to deal with balancing conflicting goals, like on one hand the wish to produce in large batches to make use of economies of scale and on the other hand to lower the inventory levels to save on the capital tied up in inventory (Axsäter, 2006). In case of perishable products, inventory control also has to balance between product-availability and waste. This thesis studies inventory control for a perishable product at several actors in the food supply chain. Therefore first a sketch of a food supply chain of a perishable product is given. Next the impact of perishability in the supply chain is briefly discussed, including issuing of product and common demand characteristics.

The supply chain of a perishable food product starts with producing the raw material at a farm. Partly the raw material (mainly fruits and vegetables) finds its way unprocessed to the consumer, in a simple package, via wholesale trade or the auction. Partly the raw material (e.g. milk, meat, fruits and vegetables) is transported to a food company where production takes place. Fig. 1.1 shows a simplified food supply chain containing the stages farm, food producer, warehouse and supermarket. In practice, the farm stage can contain several subsequent farms from breeder to fatterer of animals, or from seed producer to vegetable grower. Food production can be production of raw milk into consumer packaged milk; production of cheese; production from slaughtering to a packaged meat product; washing, cutting and packing under modified atmosphere of fresh vegetables, etc. Production steps may be executed at several subsequent production companies or locations. After production, the final products are transported to a warehouse. This might be the location of a wholesaler or a logistics service provider, possibly followed by a retail distribution centre.

Finally the products are transported to supermarkets or speciality stores to be sold to consumers.

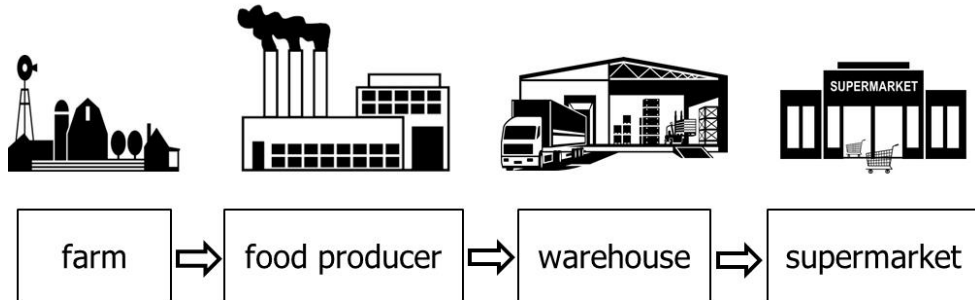


Fig. 1.1 Simplified food supply chain

Perishability has an important impact on the food supply chain. At a food producer, the product is processed and packaged in such a way that the maximum shelf life of the product is known when the product is kept in a cooled environment within a specified temperature range. After production, the product gets a best-before or use-by date on the package, reflecting this fixed maximum shelf life. This thesis focuses on perishable products with a fixed lifetime. The food producer sells the products to the customers (e.g. supermarkets) with a guaranteed remaining shelf life on the time of delivery. We define *internal shelf life* for the food producer as the maximum time span between production and distribution. To guarantee a minimum remaining shelf life at the customer, one sets a maximum on the internal shelf life. A product reaching its maximum internal shelf life means that the product cannot be sold anymore with the aimed freshness. A long enough internal shelf life allows the food producer some freedom in the production planning and distribution. Production of fast moving products usually takes place in every planning period, but planning for production of slow moving products is often more complicated. Setting up a production run costs time and money and therefore it might not be feasible to run a production for a slow moving product in every planning period. The question for a food producer is then when to plan a production run and for how many items.

Mostly, after production, produced fresh food products are distributed to the supermarket via a warehouse of a logistics service provider or a retail organisation. The limited shelf life of the products requires efficient logistics and inventory management. On the supply side, the question for a warehouse or a retail organisation is when to order and how much. A warehouse has to deliver products to many retailers. In practice it is common that large supermarkets are delivered every day, but small stores are delivered on fixed days of the week, which has to be considered in the supply of the supermarkets. The delivery of many stores has to be planned, so the question is how often and when to deliver the different types of stores. Then, the question for the retailer is how much to order, given the possible delivery days.

In determining the production or order quantity one has to consider the demand for a product and the issuing of the items from the available inventory. It is likely that the inventory of final products at the producer and at the warehouse consists of items of different ages, that is, with different production dates. The food producer and the warehouse have control over the issuing of the items. To prevent waste, the product issuing is according to a specific policy such as a First In First Out (FIFO) policy.

In supermarkets of industrialised countries, there is a wide range of fresh food products available in large quantities, leading to food waste (FAO, 2011). Consumers expect the supermarket shelves being well filled and frequent replenishment means that food products close to expiry are often ignored by consumers (FAO, 2011), causing waste. The retailer has no complete control over the withdrawal of the items. Some customers will pick the oldest items at the front of the shelf, thus picking FIFO, but other customers are determined to buy the freshest items from the back of the shelf, so they search for Last In First Out (LIFO). At the retailer often a combined LIFO – FIFO demand can be observed, thus leaving older items in stock at the expense of fresher items. It is a challenge to find the balance between well filled shelves to meet consumer demand and prevention of waste. It is important to have good insight in the factors that determine consumer demand to obtain accurate demand forecasts and to handle supply of fresh products in an appropriate way.

From practice it is known that consumer demand for perishable products varies with the seasons and the weather, as well as due to promotions at the retailer. Moreover, at the retailer a weekly demand pattern may be observed (Van Donselaar et al., 2006) showing that consumer demand per day at the end of the week is higher than in the first days of the week. This results in a non-stationary demand for products in various ways. First, at the retailer, there is a weekly (non-stationary) demand pattern which can be stationary over the weeks. In case of seasonal or weather effects, or in case of a promotion, demand can also be non-stationary over the weeks. The variations in consumer demand are transferred to the food producer. It is well known that in such cases the bullwhip effect might occur, due to lack of coordination in the supply chain (Lee et al., 1997). However, even when cooperation in the supply chain exists, there will still be a non-stationary demand at the food producer. The implications from promotions for perishable products differ from the implications for non-perishable products. Due to product perishability, the consumer cannot stash up, so demand after the promotion will recover soon. For non-perishable products, the consumer can stash up, meaning that the next buy of the product may be only after weeks or months, causing a dip in demand. For a food producer, the product perishability implies that there are hardly possibilities to produce upfront to anticipate a promotion.

Important questions in the food supply chain, at the food producer, as well as at the warehouse and the supermarket, are how much to order and how often. For perishable products, there exists a trade-off between desired product-availability and waste. A high order quantity or frequent ordering may result in high product-availability and waste. A smaller order quantity or less frequent ordering may cause out-of-stock. It is clear that for

the inventory control of a perishable product the age-distribution of the items should be considered, which is influenced by the FIFO or LIFO issuing and withdrawal of the product.

To consider all mentioned aspects in the inventory control of the actors in the food supply chain, a complex quantitative analysis is required. Modern inventory control is based on advanced and complex decision models from Operations Research which may need extensive computational effort (Axsäter, 2006). Operations Research (OR) is a field that studies the application of advanced analytical methods for making better decisions, using mathematical modelling and optimisation as well as statistical analysis. OR aims to find optimal or near-optimal solutions to complex decision-making problems (INFORMS, 2016). In this thesis, OR models are used to support inventory control in the food supply chain.

1.2 Inventory control

First the key questions of inventory control are discussed in general, without special concern for perishability. This is followed by the relevant costs to consider and a discussion on service levels. Finally, perishability in inventory literature is briefly discussed.

1.2.1 Key questions

Inventory control concerns answering three key questions for every product (Silver et al., 1998):

1. How often should the inventory status be reviewed?
2. When should an order be placed?
3. How much to order?

To review the inventory status, inventory management distinguishes two systems, a continuous review system and a periodic review system. In a continuous review system, a replenishment order is placed at any time when the inventory on hand drops below a certain reorder level s . One can order a fixed order quantity Q , or the order quantity aims to reach a so-called order-up-to level S . The order policies are called sQ or sS , respectively. In a periodic review system the inventory level is checked at regular time intervals, where R denotes the length of the review period. Periodic review is a common way of inventory control in production planning of the food industry, where production decisions mostly are made on a weekly basis. Also in retail a periodic review is applied, where order decisions are made on a daily basis.

To answer the second and third question, when and how much to order, insight in the demand pattern of the product under consideration is important. In case of stationary demand, the order decisions can be made based on the so-called Economic Order Quantity (EOQ), considering the level of uncertainty (Silver et al., 1998). This results in a steady planning over time, with fixed review periods. For products with non-stationary demand, production moments and order decisions will fluctuate in timing and quantity, especially for perishable products where smoothing of production or ordering is impossible, because the older items in stock can go to waste. The fluctuations in demand and proportional

uncertainty ask for a strategy to deal with the uncertainty. Bookbinder and Tan (1988) distinguish three strategies to come to an order policy for a certain planning horizon:

- Dynamic uncertainty strategy: the order quantity is decided at the beginning of every period in the planning horizon;
- Static-dynamic uncertainty strategy: the timing of the orders is known at the beginning of the planning horizon, the order quantity is adapted towards an order-up-to level;
- Static uncertainty strategy: timing and quantity are known at the beginning of the planning horizon.

The **dynamic uncertainty strategy** is the most studied strategy in inventory literature. The order policies are mainly variants of the RS and RsS policies. In an RS policy, every review period an order is placed up to level S . The RsS policy is applied in case of fixed setup cost for an order. The order quantity is decided every period. Only when the inventory level at the beginning of the period is below reorder level s , one orders up to level S . Otherwise there is no order. In case of stationary demand, the parameter values for R , s and S are constant in time. In case of non-stationary demand, the parameter values s_t and S_t may vary per period. This type of order policy has already been discussed in the 1960s. Karlin (1960) shows that a critical number policy is optimal for non-perishable products, where the critical numbers are a reorder level s_t and an order-up-to level S_t resulting in an (Rs_tS_t) policy. This "wait-and-see" approach in the critical number policies following the dynamic uncertainty strategy, could require an order incurring setup cost in almost every period. This might be undesirable for the production planning of a company, but in case of a large setup cost relative to the holding cost, this is also not optimal (Bookbinder and Tan, 1988). This strategy is responsive, so it may be applied in situations where service is important or where ordering every period is efficient, due to the scale of the demand. In cases where efficiency and planning is (also) important, the dynamic uncertainty strategy results in nervous planning, due to extra orders or canceled orders following the realised demand of the previous period(s) (Tunc et al. 2013). This might not be optimal from a management perspective.

In the **static-dynamic uncertainty strategy**, the timing of the orders is fixed at the beginning of the planning horizon. Beforehand is decided in which period to order or to produce, denoted by $Y = 1$ when an order is placed and $Y = 0$ if not. This results in replenishment cycles R_t of different lengths. The order quantities are determined in a dynamic way, responding to the uncertainty, by ordering up to a level S . For fluctuating demand and a varying replenishment cycle length, also the order-up-to level S_t varies per period. In this thesis, this policy is called an R_tS_t policy, or alternatively a **YS** policy, specifying more clearly that for each period is decided to order or not, and to which order-up-to level.

For a non-perishable product, Bookbinder and Tan (1988) developed a two-stage static-dynamic uncertainty model for single-stage probabilistic lot-sizing problems, where they included setup cost and service-level constraints. Tarim and Kingsman (2004) formulated a Mixed Integer Linear Programming (MILP) model for non-stationary stochastic demand for

the simultaneous determination of the number and timing of the replenishment orders. Tarim and Kingsman's approach provides an optimal solution. Others (e.g. Tempelmeier (2007), Rossi et al. (2010), Rossi et al. (2011a), Rossi et al. (2011b) and Tarim et al. (2011)) extended this approach.

A **static uncertainty strategy** can be a requirement in industry when planning ahead is important. In case of long lead time, more accurate forecasts become only known after the start of production when no changes in the production quantity can be made anymore. Adaptation of the order quantity just before realisation of demand is not possible, so the timing and production quantity have to be determined at the beginning of the planning horizon. The resulting order policy is denoted as **YQ** policy in this thesis. For non-perishable items with a non-stationary demand, Tempelmeier and Herpers (2011) formulated a Stochastic Single Item Uncapacitated Lot-Sizing Problem with a fill rate constraint. They found an optimal **YQ** solution.

The focus of this thesis is on the static-dynamic uncertainty strategy and the static uncertainty strategy, since for planning purposes it is important to fix at least the timing of production or ordering beforehand. For a food producer, it might not be feasible from cost or time perspective to run a production for a slow moving product in every planning period. For a warehouse, ordering of slow moving products might not be economically feasible in every planning period. When small supermarkets are delivered on fixed days of the week, a static-dynamic uncertainty strategy applies. These strategies have hardly been studied in inventory literature so far.

1.2.2 Relevant costs

In OR literature, most inventory control decisions are made under cost minimisation. Relevant costs for inventory decisions are procurement costs, holding costs, costs of shortage in the short run and system control costs (Axsäter, 2006). For perishable products, we also consider disposal costs of wasted items to be able to influence the level of waste, as reducing food loss and waste is a global aim set by the FAO, with a renewed initiative (FAO, 2015).

Procurement costs comprise of a *fixed ordering or setup cost* and a *variable purchasing or production cost per item*. The fixed ordering or setup cost is incurred every time an order is placed, or a production is started, independent of the order size. In case of ordering, one can think of the labour cost of the office clerk, or the transportation cost of the order. In case of production, one can think of the cost to adjust the machine for the current product or the cleaning between two production runs.

Holding costs can include a number of costs, like the cost of the capital tied up in the inventory, the cost of space and operations in the warehouse, insurance, taxes and potential spoilage or obsolescence (Axsäter, 2006). This thesis studies the inventory management of

a perishable product. Therefore we separately include disposal cost for items that reached their maximum (internal) shelf life.

Disposal costs can reflect that the items still have a salvage value where part of the variable purchasing cost can be regained, or that has to be paid to get rid of the items.

Shortage costs incur when the inventory on hand is insufficient to meet customer demand. In the short run this means that items are backlogged or that sales are lost, resulting in cost of backlogging or in loss of profit. Shortage costs are hard to determine, because shortage can lead to substitution by other products, or to procurement from a competitor and in the long run to loss of good will. An alternative way to deal with possible shortage in an inventory model is to add service level constraints to model the situation in practice (Axsäter, 2006). Section 1.2.3 discusses the concept of service levels.

System control costs are often ignored in the inventory theory literature. They include the costs of implementing an order policy, computing the parameters for the order policy and gathering the data to execute the order policy (Silver et al., 1998).

1.2.3 Service levels

There are two reasons to study a service level approach. First, food producers often have contracts with their customers, regarding delivery performance including service level and remaining shelf life. Second, stock-out penalty – or shortage cost is difficult to quantify (Minner and Transchel, 2010). Requiring a certain service level and therefore a certain level of safety stock, can be seen as a cost component accounting for shortage cost implicitly by having extra items in stock. Beforehand, it is not straightforward how high a penalty cost should be to guarantee a desired service level.

Two types of service levels used in practice are considered in this thesis: the α -service level, also called cycle service level, and the β -service level, also called fill rate. Chopra and Meindl (2010) define cycle service level as the probability of not having a stock-out in a replenishment cycle. A replenishment cycle is the time between two order deliveries. Setting a target α -service level is easier in order policy calculations than a fill rate requirement and therefore commonly used in textbooks to determine the safety stock (Silver et al., 1998; Chopra and Meindl, 2010). In practice, the α -service level can easily be evaluated afterwards. In a supermarket a stock-out can simply be observed. The fill rate indicates that a predefined percentage of the demand has to be fulfilled from stock. In practice, the realised fill rate can only be evaluated if the realised demand is known. This is the case for a food producer and warehouse, because they receive orders and know which part of the demand could not be met from stock. A supermarket does not know the amount of lost demand.

Chen and Krass (2001) define the difference between a mean and a minimal service level. A mean service level is measured over the complete time horizon, whereas a minimal service

level is measured every period or replenishment cycle. In this thesis, we study the impact of a minimal service level criterion, a minimal α -service level per period or a minimal fill rate per replenishment cycle. Food production companies often have service contracts with their retail customers requiring a certain fill rate service level. According to Chen and Krass (2001) a minimal service level criterion is preferred, when the service level requirement is due to a contractual obligation or a company policy.

Minner and Transchel (2010) determine replenishment quantities for perishable products in retail with a weekly demand pattern per day, but stationary expected demand per week, under minimal service level requirements. The combination of non-stationary demand and a service level approach in inventory models can mainly be found in literature about non-perishable products. Neale and Willems (2009) argue that non-stationary demand is very common. Therefore they developed a non-stationary supply chain inventory model, by formulating a single-stage inventory model that serves as a component of a multistage system, using service-level constraints to calculate safety stocks.

1.2.4 Perishability in inventory literature

Two categories of perishability can be found in inventory models according to Nahmias (1982): products have either a fixed lifetime or a random lifetime. Goyal and Giri (2001) and Bakker et al. (2012) distinguish two types of random lifetime, an age-dependent deterioration rate and a time- or inventory-dependent deterioration rate. Considering a random lifetime is appropriate for unprocessed products like fresh fruits and vegetables, where the moment at which deterioration occurs is unknown. A fixed lifetime is suitable for processed fresh food products, when a product gets a best-before or use-by date on the package, reflecting this fixed maximum shelf life. This thesis focuses on products with a fixed shelf life.

Nahmias (1982), Goyal and Giri (2001), Karaesmen et al. (2011) and Bakker et al. (2012) reviewed literature on inventory models for perishable products with a fixed lifetime. Almost all papers surveyed assume stationary demand, i.e. demand in successive periods is an independent identically distributed random variable. In early works e.g. Nahmias (1975) and Fries (1975) note that generally an optimal order policy for perishable products with a fixed lifetime should consider the age-distribution of the products in stock. Even when all perishable items are of the same age, base stock (*RS*) policies are not optimal, as argued by Tekin et al. (2001) and Haijema et al. (2007). Some papers, e.g. Haijema et al. (2007), Broekmeulen and Van Donselaar (2009) and Minner and Transchel (2010), assume a cyclic demand pattern, with a weekly expected demand pattern per day that is stationary per week. They assume negligible setup cost and follow a dynamic uncertainty strategy. To our knowledge, non-stationary demand which is not cyclic is hardly studied for perishable products, and also the static-dynamic uncertainty strategy for the order policy is barely investigated. The static-dynamic uncertainty strategy is appropriate for planning purposes. The above papers motivate the interest in policies for perishable products with non-stationary demand according to a static(-dynamic) uncertainty strategy.

1.3 Problem statement

The aim of this thesis is to contribute to better decision making regarding inventory control in the food supply chain from food producer to supermarket, in the trade-off between product-availability and waste. Methods are designed to generate practical order policies using commercial solvers for business rather than custom made solution procedures, where at least the timing of ordering or production is set beforehand. For the development of methods, several cases originating from practical inventory management problems have been studied.

The practical problems are finite time horizon single-product – single-echelon production/inventory control problems for a perishable product with a fixed lifetime and a service level requirement. The product under consideration has a non-stationary stochastic demand. The decision problems deal with fixed setup or ordering cost, holding cost and disposal cost for wasted items. The formulated specific theoretical Stochastic Programming (SP) problems consider the age-distribution of the items in stock and measure the reached service level in a probabilistic way.

In various practical problems, the required service level can imply an α -service level constraint, or a fill rate constraint, the issuing policy can be FIFO or combined LIFO – FIFO. When demand exceeds the inventory on hand, demand can be backlogged or alternatively is lost. For a retailer, demand of a fresh product is usually lost, whereas for a producer, backlogging might be possible. Lead time of the replenishment is zero, one period or very long.

1.4 Research opportunities

To contribute to the literature on order policies for a perishable product with non-stationary demand, several research opportunities (RO) were found and investigated in this thesis. The group of research opportunities covers inventory control for the actors in the food supply chain of produced fresh food from food producer to retailer, where order policies for a retailer are also suitable for a warehouse. The set of research opportunities also covers varying problem characteristics. High and low demand numbers are used with different demand distributions, α -service levels and fill rate constraints are investigated, FIFO issuing and combined LIFO – FIFO withdrawal is applied and excess demand is backlogged or lost. Finally lead time L is set from $L = 0$, to $L = 1$ and an approach to deal with long lead time is examined.

Inventory decisions are made under uncertainty, because the demand for a product will only be known after the reorder decision has been made. To deal with the uncertainty in demand, the inventory problems are modeled as Stochastic Programming (SP) problems. Stochastic programming is a class of mathematical programming that studies how to

incorporate uncertainty into decision problems (King and Wallace, 2012). Unfortunately, stochastic programming models are far more difficult to solve than the corresponding deterministic models. However, for certain data sets in certain models, deterministic models may produce a good solution (Kall and Wallace, 1994).

Bookbinder and Tan (1988) studied single-stage probabilistic lot-sizing problems, where they included setup cost and service-level constraints. They developed a static-dynamic uncertainty model, splitting the problem into two stages. The first stage determines when to order, the second how much to order. Tarim and Kingsman (2004) considered the Bookbinder and Tan approach as a basis for the formulation of a Mixed Integer Linear Programming (MILP) model for non-stationary stochastic demand for the simultaneous determination of the number and timing of the replenishment orders. Tarim and Kingsman's approach provides an optimal solution for a non-perishable product. In this thesis, this MILP approach is extended to a perishable product in several practical situations, each with different challenges.

The first research opportunity (RO1) is to find a practical MILP approach to determine YS parameter values for a food producer. The lead time is zero, an α -service level requirement applies and there is FIFO issuing.

The research question is:

1. Is it possible to construct practical solutions using commercial solvers for business use rather than custom made solution procedures for this problem?

This question is investigated in Chapter 2, entitled *MILP approximation generating a YS policy*.

Production/inventory models are often regarded as inherently multistage models, where each stage (period) a decision is made about the order quantity for every possible state (inventory level) in the system. When the practical problem is perceived this way, according to a dynamic uncertainty strategy, the SP model can be solved by Stochastic Dynamic Programming (SDP) (Kall and Wallace, 1994).

The second research opportunity (RO2) is to investigate whether SDP is suitable to find a solution for the practical problem described, having either an α -service level or fill rate requirement. To focus on the consequence of imposing service level constraints, the study considers a non-perishable product to avoid the complication of the age-distribution of perishable items in stock. All order policies are evaluated using either Monte Carlo simulation or, when possible, full enumeration of the realisations of demand.

The research question is:

2. Is SDP a suitable method to generate an order policy if a service level constraint applies?

This question is investigated in Chapter 3, entitled *SDP in inventory control: a $Q(X)$ policy*.

The third research opportunity (RO3) is to consider the practical problem of a food producer and compare various solution methods to generate a variety of order policies. The lead time is zero, an α -service level requirement applies and there is FIFO issuing. The MILP approximation found in Chapter 2, might not generate suitable policy parameters in all cases. The expected total cost function for an order-up-to level policy is a non-linear function in the order-up-to level. In order to obtain optimal values for the order-up to level, non-linear programming algorithms are used. The resulting Mixed Integer Non-Linear Programming (MINLP) approach is compared with the MILP approximation of Chapter 2, with a sample based approach and with order policies generated by SDP. Different solution methods by nature lead to different type of order policies.

The research question is:

3. In which situations is which policy and which solution method most suitable?

This question is investigated in Chapter 4, entitled *Comparing order policies and solution methods*.

The fourth research opportunity (RO4) is to investigate a practical situation in retail, where the reorder days are fixed and order-up-to levels are used. The lead time is one day, an α -service level requirement applies and there is FIFO or combined LIFO – FIFO withdrawal. There is a weekly demand pattern which is stationary over the weeks.

The research question is:

4. Are MILP generated order policy parameter values suitable in case of a lead time of one period and FIFO or combined LIFO – FIFO withdrawal in a rolling horizon application?

This question is investigated in Chapter 5, entitled *YS policy in retail*.

The fifth research opportunity (RO5) is to find a practical approach for a producer of a fresh food product with a long lead time to determine an order policy. Having a long lead time, the inventory levels at the time of delivery are unknown at the time the replenishment quantity has to be determined. This means that the production timing and quantity have to be determined at the beginning of the planning horizon. We consider a static uncertainty YQ policy under a cycle fill rate service level requirement and FIFO issuing.

The research question is:

5. Is it possible to generate a production plan for T periods using existing solvers and for which instances is the solution close to the optimal solution?

This question is investigated in Chapter 6, entitled *MILP approximation for a YQ policy*.

1.5 Thesis outline

This thesis includes a collection of five papers, each of them studying a research opportunity and investigating the corresponding research question. Chapters 2 to 6 are based on those papers.

Chapter 2 models the practical inventory problem of a food producer as an SP model and an MILP approximation generating a *YS* policy is presented. Chapter 3 studies the use of SDP in inventory control. In Chapter 4, order policies and solution methods are compared for the practical problem presented in Chapter 2. Chapter 5 extends the MILP approximation to a practical inventory problem in retail. Chapter 6 studies an order policy for a food producer facing a long lead time. Table 1.1 presents an overview of the problem characteristics for each research opportunity.

Table 1.1 Overview of the problem characteristics for each research opportunity and thesis chapter

Problem characteristics	RO1 – Ch2	RO2 – Ch3	RO3 – Ch4	RO4 – Ch5	RO5 – Ch6
Supply chain actor	Producer		Producer	Retailer	Producer
Perishability	Fixed shelf life	∞ lifetime	Fixed shelf life	Fixed shelf life	Fixed shelf life
Demand distribution	Normal	Uniform/Gamma	Normal	Poisson	Normal
Service level	α -service level	α -serv./ fillrate	α -service level	α -service level	fill rate
Issuing	FIFO	FIFO	FIFO	LIFO – FIFO	FIFO
Excess demand	Backlogging	Lost sales	Backlogging	Lost sales	Lost sales
Lead time	$L = 0$	$L = 0$	$L = 0$	$L = 1$	Long

1.6 Overview of papers in this thesis

Chapter 2:

Pauls-Worm, K.G.J., Hendrix, E.M.T., Haijema, R., van der Vorst, J.G.A.J., 2014. *An MILP approximation for ordering perishable products with non-stationary demand and service level constraints*. International Journal of Production Economics 157, 133-146.

Chapter 3:

Pauls-Worm, K.G.J., Hendrix, E.M.T., 2015. *SDP in Inventory Control: Non-stationary Demand and Service Level Constraints*, in: Gervasi, O., Murgante, B., Misra, S., Gavrilova, M.L., Rocha, A.M.A.C., Torre, C., Taniar, D., Apduhan, B.O. (Eds.), Computational Science and Its Applications -- ICCSA 2015. Springer International Publishing, pp. 397-412.

Chapter 4:

Pauls-Worm, K.G.J., Hendrix, E.M.T., Rossi, R., 2016. *Inventory control for a non-stationary demand perishable product: comparing policies*. Submitted to an international journal.

Chapter 1

Chapter 5:

Pauls-Worm, K.G.J., Hendrix, E.M.T., 2016. *On retail order policies for a perishable product*. Submitted to an international journal.

Chapter 6:

Pauls-Worm, K.G.J., Hendrix, E.M.T., Alcoba, A.G., Haijema, R., 2015. *Order quantities for perishable inventory control with non-stationary demand and a fill rate constraint*. International Journal of Production Economics (online 2015).

Chapter 2

MILP approximation generating a YS policy

Based on:

Pauls-Worm, K.G.J., Hendrix, E.M.T., Haijema, R., van der Vorst, J.G.A.J., 2014. *An MILP approximation for ordering perishable products with non-stationary demand and service level constraints*. International Journal of Production Economics 157, pp. 133-146.

Abstract

Chapter 2 studies the practical production planning problem of a food producer facing a non-stationary erratic demand for a perishable product with a fixed life time. In meeting the uncertain demand, the food producer uses a FIFO issuing policy. The food producer aims at meeting a certain α -service level at lowest cost. Every production run a setup cost is incurred. Moreover, the producer has to deal with unit production cost, unit holding cost and unit cost of waste. The production plan for a finite time horizon specifies in which periods to produce and how much.

We formulate this single item – single echelon production planning problem as a stochastic programming model with a chance constraint. We show that an approximate solution can be provided by an MILP model. The generated plan simultaneously specifies the periods to produce and the corresponding order-up-to levels. The order-up-to level for each period is corrected for the expected waste by explicitly considering for every period the expected age-distribution of the products in stock. The model assumes zero lead time and backlogging of shortages. The viability of the approach is illustrated by numerical experiments. Simulation shows that in 96.4% of the periods the service level requirements are met with an error tolerance of 1%.

2.1 Introduction

Food supply chains of processed fresh products generally include primary production (farmers), food processing industry, distribution centres of the producer or a retail organisation, retail stores and consumers (e.g. Van der Vorst et al. (2000)). In this chapter we study the practical production/inventory control problem faced by a food producer. After processing fresh ingredients into a final product and packing the product, the producer prints a best-before-date on the package of the product. Products can be meat, dairy products, fresh fruit juices and produced fresh meals. If the product is stored and handled under the required conditions, the product is presumed to have a fixed lifetime; the best-before-date is determined by adding a fixed number of days to the production date. In practice, a food producer often faces a non-stationary stochastic demand for his products, caused by, for instance, promotional activities of the retail organisation, or weather conditions. The producer has to decide at any given period (e.g. a week) whether to produce or not, and if so, how much to produce. This decision depends on the forecast of the demand, on the age-distribution of the items in stock and is influenced by factors such as the setup cost of a production run and the perishability of the product.

Food producers often have contracts with their customers, regarding delivery performance including service level and remaining shelf life. In order to meet these requirements and to determine production quantities, the producer has to balance product waste (as a result of too much inventory) and out-of-stock (as a result of too little inventory). Due to the perishability of the product, it is likely that the inventory of final products at the producer consists of items of different ages, that is, with different production dates. The producer sells the products to the customers (e.g. supermarkets) with a guaranteed remaining shelf life on the time of delivery. We define *internal shelf life* as the maximum time span between production and distribution. To guarantee a minimum remaining shelf life at the customer, one sets a maximum on the internal shelf life. For an internal shelf life of just 1 period, one can follow the order policy of the so-called Newsboy Problem (Silver et al., 1998) that produces every period with an order quantity that takes the perishability into account. When the internal shelf life is longer than 1 period, the order policy depends on the setup cost and holding cost and the aging of the products (Fries, 1975), (Nahmias, 1975). These studies assume out-of-stocks may happen but are penalised via a stock-out penalty or a unit shortage cost. An alternative approach to control more directly the availability of products to meet the demand, is by imposing a so-called service level constraint. Chopra and Meindl (2010) define *cycle service level* as *the probability of not having a stock-out in a replenishment cycle*. There are two reasons to study a service level approach. First, food producers often have contracts with their customers, regarding delivery performance including service level and remaining shelf life, as already mentioned. Second, stock-out penalty – or shortage cost is difficult to quantify (Minner and Transchel, 2010). Requiring a certain service level and therefore a certain level of safety stock, can be seen as a cost component accounting for shortage cost implicitly by having extra items in stock. Beforehand, it is not straightforward how high a penalty cost should be to guarantee a desired service level. When a penalty cost is set too high, one overachieves the service level

resulting in potentially unnecessary outdating. Moreover, in a non-stationary multi-period problem, the height of the penalty cost that results in precisely the desired service level depends on the period.

The demand for products of a food producer is not only stochastic, but usually is non-stationary. This leads to an irregular pattern of the decisions when to produce, and how much. Periodic review (R, S) or (R, s, S) order policies with a fixed reorder point s and order-up-to level S typically apply to stationary stochastic demand. Using such policies in case of non-stationary demand, leads to high stock levels and much waste, when S is set high to fulfil peaks in demand, or to low stock levels and many stock-outs during periods of peak demand, when S is kept low to keep overall waste low (Tunc et al., 2011). Therefore it is interesting to investigate order policies with time-dependent order-up-to levels and replenishment cycle lengths. As fixed order costs are for some food producers significant, production runs have to be planned carefully: for the coming T periods, one likes to optimise in which periods one may expect production runs; producing every R periods may be suboptimal. Uncertainty in demand leads to new production runs, while the inventory level is not zero yet. This gives items of different ages in stock. The age-distribution of the items in stock has to be monitored and should be considered in the issuing policy. A food producer has control over its issuing policy and often distributes its inventory according to First In First Out (FIFO). A fixed issuing policy, such as FIFO, is favoured in practice, because it is easy to use and keeps waste due to outdating low.

The practical problem discussed in this chapter is the finite time horizon single-product – single-echelon production/inventory control problem for a perishable product with a fixed lifetime, under a service-level constraint. The product has a non-stationary stochastic demand. The decision problem deals with a fixed setup cost for every production run and a FIFO issuing policy. We consider the age-distribution of the items in stock in a specific theoretical Stochastic Programming (SP) problem that deals with the service level as a chance constraint. The model uses zero lead time and in case of out-of-stock, demand is backlogged. The question addressed in this chapter is whether it is possible to construct practical solutions using commercial solvers for business use rather than custom made solution procedures. We approach this question with a Mixed Integer Linear Programming (MILP) model that generates approximate solutions of the problem. The approach extends previous work for non-perishable products of Bookbinder and Tan (1988) and Tarim and Kingsman (2004). An extension of their work by including the perishable nature of products is needed for at least two reasons: (1) applying a model for non-perishable products to a perishable product may result in a too optimistic view on stock levels causing desired service levels not to be met due to outdating of stock, and (2) it may even result in an infeasible solution when the time between production runs exceeds the shelf life of the product. The chapter shows that the extension towards perishable products is not a trivial one, and results in an approximate solution to the problem on hand.

This chapter is structured as follows. In Section 2.2, a literature review on the problem is presented. Section 2.3 describes the SP problem for a perishable product. In Section 2.4, a

deterministic MILP model for setting the policy parameters for a waste compensating replenishment cycle policy is formulated. In Section 2.5, the policy and the MILP model is illustrated. In Section 2.6, we present simulation results of 86 experiments for which the MILP is solved, and we investigate how well the MILP solution meets the desired service levels. The simulation results of the MILP model are compared to a Stochastic Dynamic Programming (SDP) benchmark. Chapter 2 ends with conclusions and topics for future research in Section 2.7.

2.2 Literature review

In order to construct a model for the practical problem under consideration, we review literature that deals with a combination of the key characteristics of the practical problem: perishability with a fixed lifetime, fixed setup or ordering cost, non-stationary demand, periodic review and a service-level constraint.

Nahmias (1982), Goyal and Giri (2001), Karaesmen et al. (2011) and Bakker et al. (2012) reviewed the literature on inventory models for perishable products with a fixed lifetime. Almost all papers surveyed assume stationary demand, i.e. demand in successive periods is an independent identically distributed random variable. Tekin et al. (2001) formulated an age-based control policy with a continuous review for perishable products with a fixed lifetime, under service-level constraints. The aging starts after unpacking the batch for consumption. As long as the items are packed in stock, the lifetime is virtually infinite. In early works e.g. Nahmias (1975) and Fries (1975) observe that in general an optimal order policy for perishables with a fixed life time should take the ages of the products in stock into account. Even when all perishable items are of the same age, base stock policies are not optimal, as argued by Tekin et al. (2001) and Haijema et al. (2007). Broekmeulen and Van Donselaar (2009) suggest a replenishment policy for perishable products at a retailer, which takes the quantity and the age of the items in inventory into account. They assume negligible fixed ordering cost. The demand is assumed to be stochastic, with a weekly demand pattern per day, but stationary expected demand per week. They apply the same safety stock for each weekday. Haijema et al. (2007) developed an optimal policy for the periodic production and inventory of blood platelets. They combine two types of demand, each of which requires a different issuing policy. The demand distributions they consider have a weekly demand pattern per day, but are stationary across weeks. In (Haijema et al., 2009) the approach is extended for non-stationary demand considering holidays and other events. Any fixed production cost is neglected. In (Haijema, 2013) fixed order cost are studied and a new class of order policies is presented. In none of these papers service-level constraints are included. Minner and Transchel (2010) present a numerical approach to determine replenishment quantities for perishable products in retail dynamically, using a weekly demand pattern. They consider service-level constraints varying for different intra-period time points and for different periods. Fixed ordering cost is assumed to be negligible.

In our investigation, the combination of non-stationary demand and a service level approach in inventory models was mainly found in literature about non-perishable products. Neale and

Willems (2009) argue that non-stationary demand is very common. Therefore they developed a non-stationary supply chain inventory model, by formulating a single-stage inventory model that serves as a component of a multistage system, using service-level constraints to calculate safety stocks. The model is based on (Graves and Willems, 2000) and closely related to (Graves and Willems, 2008). Every stage has a base-stock policy with a review period of 1 time unit. The base-stock level is an order-up-to level to cover demand in upcoming periods. The safety stock is calculated as a function of demand over the preceding periods. In the multi-stage system Neale and Willems (2009) minimise the total holding cost of the safety stock in all stages and periods. They do not consider setup cost, which is an important cost component in practice. Bookbinder and Tan (1988) studied single-stage probabilistic lot-sizing problems, where they included setup cost and service-level constraints. They developed a “static-dynamic” uncertainty model, splitting the problem in two stages. The first stage determines when to order, the second how much to order. Tarim and Kingsman (2004) considered the Bookbinder and Tan approach as a basis for the formulation of a mixed integer programming model for non-stationary stochastic demand for the simultaneous determination of the number and timing of the replenishment orders. In contrast to Bookbinder and Tan’s heuristic approach, Tarim and Kingsman’s approach provides an optimal solution. Several extensions of Tarim and Kingsman’s model exist. Rossi et al. (2011b) and Tarim et al. (2011) proposed efficient and complete special purpose algorithms. Tempelmeier (2007) used Tarim and Kingsman’s model as a basis to formulate different types of service-level constraints. Rossi et al. (2010) and Rossi et al. (2011a) incorporated a stochastic delivery lead time and developed both complete and fast heuristic approaches. Tempelmeier (2011) incorporated supplier capacity constraints. Pujawan and Silver (2008) proposed a novel and effective heuristic approach. However, to the best knowledge of the authors, no paper deals with all aspects of the practical planning problem under consideration: the combination of perishability with a fixed lifetime, fixed setup or ordering cost, non-stationary demand and a service level approach. In this chapter we extend the model of Tarim and Kingsman towards a model that includes non-stationary stochastic demand for a perishable product under a FIFO issuing policy.

2.3 SP model for a perishable product

The problem of determining a production plan for a perishable product under non-stationary stochastic demand consists of deciding when to produce and how much to produce for a finite time horizon of T periods, such that the expected total costs are minimised. Periods can be hours, days, weeks or months, whatever is applicable in the practical situation. We adopt a minimum service-level criterion for meeting customer demand. Consider a single-product – single-echelon model where the product has a fixed maximum integer (internal) shelf life $M \geq 2$ periods. A replenishment arrives instantaneously at the beginning of a period, i.e. lead time is zero. Demand d_t is a non-stationary independent stochastic process with probability density function $g_t(\cdot)$ and cumulative distribution function $G_t(\cdot)$. Demand is never negative; food cannot be returned due to food safety regulations.

We consider a FIFO issuing policy in which the first produced items are issued first. Let the ages be indexed by $b = 1, \dots, M$. Variable I_{bt} denotes the inventory level of items with age b at the end of period t . Items that are delivered at the beginning of period t have age $b = 1$ at the end of period t . Items of age M at the end of a period are not carried over to the next period, because they are out-dated; inventory I_{Mt} of age M at the end of period t is considered waste. Demand that cannot be fulfilled in one period is backlogged in the next period. Further costs are a fixed setup cost k for every production run and a variable production cost c per item produced. We assume that k and c are independent of the production period, but the model can be generalised with a period-dependent setup cost and production cost. For items that are carried over from one period to the next, a holding cost h per item is incurred. There is a cost w per item of waste, on top of the unit production cost c . The case $w > 0$, describes a situation with additional cost to discard the wasted items. Situation $w < 0$ reflects that the wasted items still have a salvage value of $-w$. All costs remain constant within the time horizon. For convenience and without loss of generality, the initial inventory level is set to zero. An overview of the used symbols is presented in Appendix 2.A. The resulting problem can be formulated as a stochastic programming model:

$$\text{Min } E(TC) = \int_{d_1} \dots \int_{d_T} \sum_{t=1}^T \left(h \max\{I_{1t}, 0\} + h \sum_{b=2}^{M-1} I_{bt} + f(Q_t) + w I_{Mt} \right) g_1(d_1) \dots g_T(d_T) dd_1 \dots dd_T \quad (1)$$

subject to

$$f(Q_t) = \begin{cases} k + cQ_t & \text{if } Q_t > 0 \\ 0 & \text{otherwise} \end{cases} \quad t = 1, \dots, T \quad (2)$$

$$\sum_{b=1}^M I_{bt} = \sum_{b=1}^{M-1} I_{b,t-1} + Q_t - d_t \quad t = 1, \dots, T \quad (3)$$

$$\mathbb{P}\left(\sum_{b=1}^M I_{bt} \geq 0\right) \geq \alpha \quad t = 1, \dots, T \quad (4)$$

$$I_{bt} = \max\left\{I_{b-1,t-1} - \max\left\{d_t - \sum_{j=b}^{M-1} I_{j,t-1}, 0\right\}, 0\right\} \quad t = 1, \dots, T; \quad b = 2, \dots, M \quad (5)$$

$$I_{1t} = Q_t - \max\left\{d_t - \sum_{b=1}^{M-1} I_{b,t-1}, 0\right\} \quad t = 1, \dots, T \quad (6)$$

$$I_{b0} = 0 \quad b = 1, \dots, M \quad (7)$$

$$I_{bt} \geq 0 \quad t = 1, \dots, T; \quad b = 2, \dots, M \quad (8)$$

$$I_{1t} \in \mathbb{R}, \quad Q_t \geq 0 \quad t = 1, \dots, T \quad (9)$$

The objective function Eq. (1) of the model minimises the expected total costs, comprising fixed setup cost for every production run, holding cost over every item in stock, unit production cost and cost of wasted items. In Eq. (3), the inventory levels of all ages are

balanced. Items of age M cannot be used in the next period, so period t starts with the inventory levels at the end of period $t - 1$ of ages $b = 1, \dots, M - 1$. The inventory at the end of period t equals the starting inventory increased by an amount Q_t that is produced in period t minus the demand in period t . Service-level constraint Eq. (4) states that the inventory levels of all ages together at the end of period t should be nonnegative with probability α . This type of service level is known as α -service level. Eq. (5) and (6) are the FIFO constraints. They make sure that demand is fulfilled first by the oldest items in stock and then successively by the younger items. Possible shortages only occur for the youngest items (Eq.(6)). Notice that adding up all equations of (5) and (6) results in Eq. (3). The starting inventory level of all ages is 0 (Eq. (7)), and the inventory levels of all ages in all other periods are nonnegative (Eq. (8)) except for the inventory level of age 1 in all periods, which can be negative when stock is too small to fulfil demand as defined in Eq. (9). Compared to the model of Bookbinder and Tan (1988), the SP model also considers FIFO constraints Eq. (5) and (6), and included the age of the items to the variable for the inventory level.

2.4 Deterministic MILP approximation

We show how a deterministic MILP model can generate a waste-compensating replenishment cycle policy as an approximate solution of the SP model for a perishable product. Therefore, we first discuss in Section 2.4.1 the structure of the replenishment cycle policy. Sections 2.4.2 to 2.4.4 present the following ingredients of the MILP model: the objective function Eq. (1), the service-level constraint Eq. (4) and the FIFO constraints Eq. (5) and (6). In Section 2.4.5, the complete model is presented.

2.4.1 The waste-compensating replenishment cycle policy

In the SP model, the structure of the policy is not explicit: the order quantity Q_t can be regarded as a stochastic variable as it depends on the inventory levels which in turn depend on the stochastic demand. The MILP approximation that we derive in the next section, optimises for each period of the horizon two order parameters: (1) Y_t that indicates whether a production order is placed or not in period t , and (2) a fixed period-dependent order-up-to level or target starting inventory level S_t . Compared to the value of S_t in non-perishable inventory systems, in perishable inventory systems, S_t should compensate for the expected waste taking into account a desired service level. Therefore, we call the policy a “waste-compensating” replenishment cycle policy. At the beginning of the planning horizon, we simultaneously determine in which periods to produce defined by Y_t and the associated order-up-to levels S_t which aim at fulfilling the prescribed service level. The order-up-to level S_t for period t , should not only cover demand in the replenishment cycle up to the moment the next production run will take place, but should also compensate for (expected) waste throughout the replenishment cycle. Therefore, the inventory level at the end of period $t - 1$ and the age-distribution of the inventory should be monitored. The production quantity Q_t is determined by the order-up-to level minus the inventory on hand at the end of in period

$t - 1$. In case the inventory on hand exceeds the order-up-to level, the excess stock will be carried forward. The order quantity is defined by

$$Q_t = Y_t \cdot \max \left\{ S_t - \sum_{b=1}^{M-1} I_{bt}, 0 \right\} \quad t = 1, \dots, T \quad (10)$$

Notice again that the quantity itself is a random variable due to the stochastic nature of demand, whereas the parameters S_t and Y_t are not. Replenishment cycles can vary in length from 1 to M periods.

2.4.2 Objective function

Consider the objective function (1). The holding cost is given by

$$\int \dots \int_{d_1} \sum_{d_T} \left(h \max \{ I_{1t}, 0 \} + h \sum_{b=2}^{M-1} I_{bt} \right) g_1(d_1) \dots g_T(d_T) dd_1 \dots dd_T \quad (11)$$

where no holding cost is paid over negative inventory. The service level α is usually chosen such that the probability of out-of-stock $1 - \alpha$ is small. The approximation assumes that the occurrence and amount of shortage is small enough to be neglected in the calculation of the holding cost (Bookbinder and Tan, 1988). That gives the objective function

$$\text{Min } E(TC) = \int \dots \int_{d_1} \sum_{d_T} \left(h \sum_{b=1}^{M-1} I_{bt} + f(Q_t) + wI_{Mt} \right) g_1(d_1) \dots g_T(d_T) dd_1 \dots dd_T. \quad (12)$$

In the waste-compensating replenishment cycle policy defined by Eq. (10), the binary variable Y_t is fixed, leaving Q_t as stochastic variable. Then Eq. (12) is equivalent to

$$\text{Min } E(TC) = \sum_{t=1}^T \left\{ kY_t + h \sum_{b=1}^{M-1} E(I_{bt}) + cE(Q_t) + wE(I_{Mt}) \right\} \quad (13)$$

Eq. (13) is the objective function of the deterministic MILP model. Expressions for the expectation of $\sum_{b=1}^M I_{bt}$ and Q_t as function of the expected demand d_t are straightforward.

Expected values for the separate variables I_{bt} are more complicated as we will specify when considering the FIFO constraints.

2.4.3 Service level constraint

The variable S_t denotes the order-up-to level or target starting inventory level at the beginning of period t . When no order is placed, $Q_t = 0$ and S_t equals the ending inventory level of period $t - 1$. When an order is placed, S_t is the order-up-to level and Eq. (3) becomes

$$\sum_{b=1}^M I_{bt} = S_t - d_t \quad t = 1, \dots, T \quad (14)$$

Eq. (4) requires the inventory level at the end of every period to be nonnegative with a probability of service level α . Using Eq. (14), Eq. (4) can be rewritten as

$$P(d_t \leq S_t) \geq \alpha \quad t = 1, \dots, T \quad (15)$$

specifying that the order-up-to level or starting inventory level of every period should be greater than the demand with a probability higher than the service level. Now consider period t , when the last order prior to period t took place in period $t-j+1$, to fulfil demand of j periods. The next order is in period $t+1$, with $j \in \{1, \dots, M\}$. So, for example, consider $M \geq 3$. When $j = 3$ then $Y_{t-2} = 1$, $Y_{t-1} = 0$, $Y_t = 0$, $Y_{t+1} = 1$. When $j = 2$ then $Y_{t-1} = 1$, $Y_t = 0$, $Y_{t+1} = 1$. Let $G_{t-j+1,t}(\cdot)$ be the cumulative probability distribution function of $d_{t-j+1} + d_{t-j+2} + \dots + d_t$. To meet the desired service level we need

$$P\left(\sum_{n=t-j+1}^t d_n \leq S_{t-j+1}\right) \geq \alpha \quad t = 1, \dots, T \quad (16)$$

which implies

$$G_{t-j+1,t}(S_{t-j+1}) \geq \alpha \quad \text{or} \quad S_{t-j+1} \geq G_{t-j+1,t}^{-1}(\alpha) \quad t = 1, \dots, T \quad (17)$$

such that

$$\sum_{b=1}^M I_{bt} \geq G_{t-j+1,t}^{-1}(\alpha) - \sum_{n=t-j+1}^t d_n \quad t = 1, \dots, T. \quad (18)$$

Eq. (18) specifies that the inventory level at the end of period t depends on the convolution of the probability distribution of the demand over the previous periods, since the last order took place. The inventory level at the end of period t may consist of items of different ages, including items of age M that cannot be used in period $t+1$. The expected value of the right-hand-side of Eq. (18) represents the safety stock at the α -service level. The required safety stock is known, given the period the order takes place. In this model, a finite planning horizon of T periods is considered. Therefore, $G_{t-j+1,t}^{-1}(\alpha)$ can be calculated in advance for all relevant combinations of j and t . Let Z_{tj} be a binary variable that is equal to 1 if the most recent order prior to period t was in period $t-j+1$, then holds

$$\sum_{b=1}^M I_{bt} \geq \sum_{j=1}^M \left(G_{t-j+1,t}^{-1}(\alpha) - \sum_{n=t-j+1}^t d_n \right) \cdot Z_{tj} \quad t = 1, \dots, T. \quad (19)$$

As only one order period is the most recent order period prior to period t ,

$$\sum_{j=1}^M Z_{tj} = 1 \quad t = 1, \dots, T \quad (20)$$

and

$$Z_{tj} \geq Y_{t-j+1} - \sum_{n=t-j+2}^t Y_n \quad t = 1, \dots, T ; j = 1, \dots, M. \quad (21)$$

Eq. (21) specifies that if $Z_{tj} = 1$, then the most recent order prior to period t was in period $t-j+1$, so $Y_{t-j+1} = 1$, and Y_{t-j+2} to Y_t should all be 0. Otherwise, Z_{tj} will be 0. Eq. (19) can be interpreted as the calculation of the safety stock needed to fulfil demand from periods $t-j+1$ to t , when there is an order in period $t-j+1$. The safety stock, or the inventory level at the end of period t may consist of items of different ages, including items of age M ,

which can be used in period t , but not in period $t + 1$. These items are considered waste at the end of period t . Using expected values, Eq. (19) becomes

$$\sum_{b=1}^M E(I_{bt}) \geq \sum_{j=1}^M \left(G_{t-j+1,t}^{-1}(\alpha) - \sum_{n=t-j+1}^t E(d_n) \right) \cdot Z_{tj} \quad t = 1, \dots, T \quad (22)$$

2.4.4 FIFO constraints

Constraints (5) and (6) make sure that items are issued according to a FIFO policy. To obtain insight in them, consider $M = 3$, with different values for the index b . In this setting Eq. (5) implies equations (23) and (24), and Eq. (6) becomes Eq. (25):

$$I_{3t} = \max\{I_{2,t-1} - \max\{d_t - 0, 0\}, 0\} = \max\{I_{2,t-1} - d_t, 0\} \quad t = 1, \dots, T \quad (23)$$

$$I_{2t} = \max\{I_{1,t-1} - \max\{d_t - I_{2,t-1}, 0\}, 0\} \quad t = 1, \dots, T \quad (24)$$

$$I_{1t} = Q_t - \max\{d_t - I_{1,t-1} - I_{2,t-1}, 0\} \quad t = 1, \dots, T \quad (25)$$

To construct a deterministic MILP model, the expected values of these constraints are needed. Since the function $\max\{\cdot\}$ is a convex function, Jensen Inequality (Mood et al., 1974) applies. The MILP model uses the following equalities to approximate the SP model.

$$E(I_{3t}) = \max\{E(I_{2,t-1}) - E(d_t), 0\} \quad t = 1, \dots, T \quad (26)$$

$$E(I_{2t}) = \max\{E(I_{1,t-1}) - \max\{E(d_t) - E(I_{2,t-1}), 0\}, 0\} \quad t = 1, \dots, T \quad (27)$$

$$E(I_{1t}) = E(Q_t) - \max\{E(d_t) - E(I_{1,t-1}) - E(I_{2,t-1}), 0\} \quad t = 1, \dots, T \quad (28)$$

According to Jensen Inequality the expected waste in Eq. (26) is underestimated and the expected inventory level of the freshest items in Eq. (28) is overestimated. Due to the nested function $\max\{\cdot\}$ in Eq. (27), Jensen Inequality does not apply and the approximation could be an under- or overestimation.

Consider Eq. (27). If $E(I_{2,t-1}) - E(d_t) \geq 0$ then $E(I_{2,t-1}) - E(d_t) = E(I_{3t})$. If $E(I_{2,t-1}) - E(d_t) \leq 0$ then $E(d_t) - E(I_{2,t-1}) \geq 0$. The value of $E(d_t) - E(I_{2,t-1}) \geq 0$ can be seen as the residual demand for the oldest items. This amount has to be fulfilled by fresher items as can be seen in Eq. (27). Let the auxiliary variable $E(X_{bt})$ denote the residual demand for items of age b with $b = 1, \dots, M - 1$ in period t . If $E(X_{bt})$ has a positive value, then fresher inventory is used to fulfil demand:

$$E(X_{2t}) = E(d_t) - E(I_{2,t-1}) \geq 0. \quad (29)$$

Using $E(X_{2t})$, Eq. (26) becomes

$$I_{2,t-1} - d_t = I_{3t} - X_{2t} \quad (30)$$

and Eq. (27) becomes

$$E(I_{2t}) = \max\{E(I_{1,t-1}) - E(X_{2t}), 0\} \quad \text{or} \quad (31)$$

$$E(I_{1,t-1}) - E(X_{2t}) = E(I_{2t}) - E(X_{1t}).$$

Finally, Eq. (28) can be formulated as

$$E(Q_t) - E(X_{1t}) = E(I_{1t}), \quad (32)$$

where $E(I_{1t}) \geq 0$, because the deterministic model assumes there are no out-of-stocks. For $M = 3$, Eq. (5) and (6) are equivalent to:

$$\begin{cases} E(I_{2,t-1}) - E(d_t) = E(I_{3t}) - E(X_{2t}) \\ E(I_{1,t-1}) - E(X_{2t}) = E(I_{2t}) - E(X_{1t}) \\ E(Q_t) - E(X_{1t}) = E(I_{1t}) \end{cases} \quad t = 1, \dots, T \quad (33)$$

This set of equations handles the age-distribution of the items in stock. Adding up the equations of Eq. (33) results into

$$\sum_{b=1}^M E(I_{bt}) = \sum_{b=1}^{M-1} E(I_{b,t-1}) + E(Q_t) - E(d_t) \quad t = 1, \dots, T. \quad (34)$$

Eq. (34) is equivalent to Eq. (3) of the SP model. In Section 2.4.5, the complete set of FIFO constraints is written in Eq. (42) to (46).

2.4.5 MILP model for a perishable product

The complete deterministic MILP model to generate approximate policies of the SP model is presented below.

$$\text{Min } E(TC) = \sum_{t=1}^T \left\{ kY_t + h \sum_{b=1}^{M-1} E(I_{bt}) + cE(Q_t) + wE(I_{Mt}) \right\} \quad (35)$$

$$\sum_{b=1}^M E(I_{bt}) = S_t - E(d_t) \quad t = 1, \dots, T \quad (36)$$

$$E(Q_t) \leq \mathcal{M}Y_t \quad t = 1, \dots, T \quad (37)$$

\mathcal{M} is a sufficiently large number, for instance $\mathcal{M} = \sum_{t=1}^T E(d_t)$. This upper bound is high

enough, as the order quantity in the deterministic MILP model will not exceed the expected demand over all T periods, because shelf life $M < T$.

$$S_t \geq \sum_{j=1}^M \left(G_{t-j+1,t}^{-1}(\alpha) - \sum_{n=t-j+1}^t E(d_n) \right) \cdot Z_{tj} + E(d_t) \quad t = 1, \dots, T \quad (38)$$

Eq. (38) is a reformulation of Eq. (22) using the variable S_b instead of $E(I_{bt})$, to obtain the desired order-up-to levels to meet the α -service level requirement.

$$\sum_{j=1}^M Z_{tj} = 1 \quad t = 1, \dots, T \quad (39)$$

$$Z_{tj} \geq Y_{t-j+1} - \sum_{n=t-j+2}^t Y_n \quad t = 1, \dots, T ; j = 1, \dots, M \quad (40)$$

$$E(Q_t) = S_t - \sum_{b=1}^{M-1} E(I_{b,t-1}) \quad t = 1, \dots, T. \quad (41)$$

Eq. (41) specifies the expected required production quantity. Waste is considered by leaving out items $E(I_{M,t-1})$, because they cannot be used in period t . Let the auxiliary variable $E(X_{bt})$ denote the residual demand for items of age b with $b = 1, \dots, M - 1$ in period t . If $E(X_{bt})$ has a positive value, then fresher inventory is used to fulfil demand.

$$E(I_{M-1,t-1}) - E(d_t) = E(I_{Mt}) - E(X_{M-1,t}) \quad t = 1, \dots, T \quad (42)$$

$$E(I_{b,t-1}) - E(X_{b+1,t}) = E(I_{b+1,t}) - E(X_{bt}) \quad t = 1, \dots, T; \quad b = 1, \dots, M - 2 \quad (43)$$

$$E(Q_t) - E(X_{1t}) = E(I_{1t}) \quad t = 1, \dots, T \quad (44)$$

Eq. (42), (43) and (44) keep track of the age-distribution of the items in stock, under a FIFO-issuing policy. Eq. (42) imposes the oldest inventory to be used first to fulfil demand. What is left over has the maximum shelf life and will become waste, or there will be a residual demand for the oldest items. In the latter case, Eq. (43) is appropriate. The residual demand has to be fulfilled by items of intermediate ages, until the demand is fulfilled by the freshest items that are produced in the current period, according to Eq. (44). The right-hand-sides of equations (42) and (43) can each contain at most one variable with a positive value. The other variable needs to have a value of 0. Equations (45), and (46) impose that, using the binary variable BX_{bt} ,

$$\mathcal{M} \cdot BX_{bt} \geq E(X_{bt}) \quad t = 1, \dots, T; \quad b = 1, \dots, M - 1 \quad (45)$$

$$\mathcal{M} \cdot (1 - BX_{bt}) \geq E(I_{b+1,t}) \quad t = 1, \dots, T; \quad b = 1, \dots, M - 1 \quad (46)$$

$$E(I_{b0}) = 0 \quad b = 1, \dots, M \quad (47)$$

$$E(I_{bt}), E(Q_t), S_t \geq 0 \quad t = 1, \dots, T; \quad b = 1, \dots, M \quad (48)$$

$$E(X_{bt}) \geq 0 \quad t = 1, \dots, T; \quad b = 1, \dots, M - 1 \quad (49)$$

$$Y_t, Z_{tj} \in \{0, 1\} \quad t = 1, \dots, T; \quad j = 1, \dots, M \quad (50)$$

$$BX_{bt} \in \{0, 1\} \quad t = 1, \dots, T; \quad b = 1, \dots, M - 1 \quad (51)$$

The starting inventory is zero (Eq.(47)). Eq. (48) to (51) are definition constraints. The required safety stocks are part of the inventory level. This model is an extension of the MILP model formulated by Tarim and Kingsman (2004), considering the age-distribution of the items in stock and the FIFO constraints. Their model provides an optimal solution for an SP model for non-perishables. Specifically, the addition of the FIFO constraints makes that the MILP model for perishable products generates approximate solutions for the SP model.

2.5 Numerical illustration of the MILP model

In Section 2.5.1, we illustrate the waste-compensating replenishment cycle policy of the model with a numerical example. The chosen parameter values are extreme to demonstrate the behaviour of the model. In Section 2.5.2, we relax the FIFO constraints for the same numerical example.

2.5.1 Numerical illustration of the MILP model with FIFO issuance

A numerical example illustrates the waste-compensating replenishment cycle policy of the model for a product with a fixed (internal) shelf life of 3 periods. The chosen parameter values are somewhat extreme to emphasize the particular behaviour of the model. In Section 2.6, the model is tested for another 86 different, more realistic, problem settings and compared to an optimal policy obtained by Stochastic Dynamic Programming.

In the example, the demand in each period is normally distributed, with a Coefficient of Variation (CV) of 0.333. Demand forecasts $E(d_t)$ are given in Table 2.1. The fixed setup cost is set to $k = 3000$, the unit production cost to $c = 2$ and the holding cost to $h = 1$. The cost of waste or markdown of the product w is 4. We use a service level of $\alpha = 95\%$.

Table 2.1 Forecasts and standard deviations of demand with a constant CV = 0.333

Period t	1	2	3	4	5	6	7	8	9	10	11	12
$E(d_t)$	1900	950	40	80	30	150	800	950	1100	350	150	700
$St.dev.(d_t)$	632.7	316.4	13.32	26.64	9.99	49.95	266.4	316.4	366.3	116.6	49.95	233.1

The safety stocks in Eq. (38) to meet a 95% service level are given in Table 2.2. For example, the safety stock at the end of period $t = 3$ is 521 (highlighted in Table 2.2) when the most recent order prior to period 3 was in period $t - j + 1 = 3 - 2 + 1 = 2$, for $j = 2$ periods.

Table 2.2 Safety stocks when the most recent order prior to period t was in period $t - j + 1$: ordering for j periods

j	t	1	2	3	4	5	6	7	8	9	10	11	12
1		1041	521	22	44	17	83	439	521	603	192	83	384
2			1164	521	49	47	84	446	681	797	633	209	393
3				1164	523	52	95	447	686	909	819	638	437

The MILP model provides the policy given in Table 2.3. Orders occur in periods 1, 2, 4, 7, 9, 10 and 12, seven times. The order-up-to level of period 2 is equal to 1511. This amount is to fulfil demand of periods 2 and 3 and the safety stock at $t = 3$ and $j = 2$ according to Table 2.2. The actual amount ordered is equal to the order-up-to level minus the inventory at the end of period 1. To fulfil demand in period 2, the one-period-old items of period 1 are used. In period 4 a production run takes place, to fulfil demand of periods 4, 5 and 6 and the safety stock for these periods. So $80 + 30 + 150 + 95 = 355$ items are required, and one would expect an order-up-to level of 355. We call this the basic order-up-to level. Instead, $S_4 = 745$ in Table 2.3. Note that the demand of period 4 is fulfilled by the two-periods-old items of period 3. Afterwards, there is 390 items expected waste, which cannot be used in periods 5 and 6. The order-up-to level of period 4 is corrected for the amount of waste: $355 + 390 = 745$, so the order-up-to level is waste-compensating. As illustrated, the MILP model determines order-up-to levels taking into account the expected age-distribution of the inventory rather than the actual inventory level.

Table 2.3 Order policy and model output for the example problem, CV = 0.333

t	1	2	3	4	5	6	7	8	9	10	11	12
$E(d_t)$	1900	950	40	80	30	150	800	950	1100	350	150	700
S_t	2941	1511	561	745	275	245	2431	1631	1703	709	359	1084
$E(Q_t)$	2941	470	0	275	0	0	2431	0	1022	106	0	978
$\sum_b E(I_{bt})$	1041	561	521	665	245	95	1631	681	603	359	209	384
$E(I_{1t})$	1041	470	0	275	0	0	1631	0	603	106	0	384
$E(I_{2t})$	0	91	470	0	245	0	0	681	0	253	106	0
$E(waste_t)$	0	0	51	390	0	95	0	0	0	0	103	0

2.5.2 FIFO relaxation

In the practical decision problem a FIFO issuing policy is used. A different approach to determine a production plan is to use no predetermined issuing policy. Therefore, the strict FIFO constraints (42) to (46) are replaced by the inventory balance constraints (52) to (54).

$$\sum_{b=1}^{M-1} E(I_{b,t-1}) - E(d_t) + E(Q_t) = \sum_{b=1}^M E(I_{b,t}) \quad t = 1, \dots, T \quad (52)$$

$$E(I_{b,t-1}) \geq E(I_{b+1,t}) \quad t = 1, \dots, T; \quad b = 1, \dots, M-1 \quad (53)$$

$$E(Q_t) \geq E(I_{1,t}) \quad t = 1, \dots, T \quad (54)$$

The FIFO constraints in the model explicitly take care that the oldest items are issued first. If demand is less than the inventory in stock of items of age two at the end of the previous period, there will be waste. This can be seen in Table 2.3. E.g. at the end of period 2 $E(I_{2t}) = 91$. Expected demand in period 3 is 40, so expected waste at the end of period 3 is 51. In the same way, the expected waste at the end of period 4 is 390.

If we relax the FIFO constraints to inventory balance constraints, the issuing will be done in a cost-optimal way. Application of the relaxation to the numerical example of Section 2.5.1, results into the policy shown in Table 2.4. It can be observed that the expected demand of period 2 of 950 items is fulfilled by 390 of the 470 fresh items (Q_2), and by 560 items of the items of age one at the end of the previous period, while more items of age one are still available. The items of age one at the end of period 2 are now of age two and they are used to fulfil demand of period 3 and the rest goes to waste. In the remaining of the time horizon a FIFO issuing is applied.

Due to the relaxation of the FIFO constraints, the expected total costs of this production plan are reduced to 45968 compared to those of the FIFO production plan of 46358. For this specific instance, the amount of waste is the same, but the timing is different, resulting in lower holding cost. In the relaxation there is no clear issuing rule: sometimes FIFO, sometimes mixed ages. In practice of a production location, a clear issuing rule is necessary, so the FIFO relaxation is not suitable.

Table 2.4 Order policy and model output for the example problem, $CV = 0.333$, no issuing policy

t	1	2	3	4	5	6	7	8	9	10	11	12
$E(d_t)$	1900	950	40	80	30	150	800	950	1100	350	150	700
S_t	2941	1511	561	355	275	245	2431	1631	1703	709	359	1084
$E(Q_t)$	2941	470	0	275	0	0	2431	0	1022	106	0	978
$\Sigma_b E(I_{bt})$	1041	561	521	275	245	95	1631	681	603	359	209	384
$E(I_{1t})$	1041	80	0	275	0	0	1631	0	603	106	0	384
$E(I_{2t})$	0	481	80	0	245	0	0	681	0	253	106	0
$E(waste_t)$	0	0	441	0	0	95	0	0	0	0	103	0

2.6 Results of the MILP model

This section investigates the behaviour of the model for different parameter values and different demand patterns. Section 2.6.1 describes the design of experiments. In Section 2.6.2, Monte Carlo simulation measures how well the MILP policies meet the service level requirements. We study the influence of varying setup cost, cost of waste, service level, coefficient of variation, maximum shelf life and demand pattern, on expected total costs, expected waste and the production plan. We also compare the performance of the MILP approximation with that of an SDP model. In Section 2.6.3, a benchmark is presented of the simulation results of the MILP policy with the SDP policy for all experiments with $M = 3$.

With respect to the fixed setup cost k there are three situations:

1. In the situation where the setup cost k is low and we order every period, the levels S_t are given by the basic order-up-to levels composed from the safety stocks such as shown in Table 2.2 added to the expected demand. This policy is exactly generated by the MILP model. There is no waste during the replenishment cycle, only at the end of the cycle. The service levels are fulfilled. SDP generates almost the same order policy. Table 2.5 shows the results for the base case which will be defined in Section 2.6.1, but with $k = 0$.
2. If k is high and we order every M periods, the starting inventory levels of all replenishment cycles are zero and the value of the order-up-to level S_t based on the cumulative distribution function is exactly the value that provides the service level. There is no waste during the replenishment cycle, only at the end of the cycle. SDP generates a more strict order policy with higher service levels at higher costs, because SDP starts a production run more often: on average 4.36 times instead of 4 times. Table 2.6 shows the result for the base case, but with $k = 4000$.
3. If k is intermediate, the replenishment cycles are not always of the same length. This is the challenging situation. The starting inventory levels differ and waste can occur during a replenishment cycle. In the design of experiments, we focus on this situation, to find when the service level is not met.

Table 2.5 Setup cost $k = 0$: MILP production plan and MILP simulation results versus SDP results (10,000 runs)*

	T	1	2	3	4	5	6	7	8	9	10	11	12
Data	$E(d_t)$	800	950	200	900	800	150	650	800	900	300	150	600
	S_t	1129	1341	391	1271	1129	329	918	1129	1271	424	212	847
	$E(Q_t)$	1129	1012	0	1080	758	0	739	861	942	53	88	785
MILP	$E(I_{1t})$	329	391	0	371	329	0	268	329	371	53	62	247
	$E(I_{2t})$	0	0	191	0	0	179	0	0	0	71	0	0
	$E(waste_t)$	0	0	0	0	0	0	0	0	0	0	0	0
Sim	α -serv.	95.0	95.4	100.0	95.4	94.8	100.0	95.2	95.4	95.1	96.8	96.6	95.2
	AvQ	1129	1011	49	1030	760	34	704	864	942	118	66	804
	$\# orders$	1.00	1.00	0.33	1.00	1.00	0.27	1.00	1.00	1.00	0.59	0.66	1.00
	$Inv1$	329	389	21	369	321	13	265	327	367	54	32	248
SDP	$Inv2$	0	2	220	0	10	200	0	2	4	133	11	0
	$Waste$	0	0	0	1	0	3	3	0	0	0	60	0
	α -serv.	95.0	95.3	98.1	95.3	94.8	98.4	95.2	95.4	95.0	96.8	96.6	95.1

* MILP $E(TC) = 16489.5$; Simulated MILP Average $TC = 16749$; SDP Average $TC = 16693$.**Table 2.6** Setup cost $k = 4000$: MILP production plan and MILP sim. results versus SDP results (10,000 runs)*

	T	1	2	3	4	5	6	7	8	9	10	11	12
Data	$E(d_t)$	800	950	200	900	800	150	650	800	900	300	150	600
	S_t	2468	1668	718	2350	1450	650	2913	2263	1463	1333	1033	883
	$E(Q_t)$	2468	0	0	2350	0	0	2913	0	0	1333	0	0
MILP	$E(I_{1t})$	1668	0	0	1450	0	0	2263	0	0	1033	0	0
	$E(I_{2t})$	0	718	0	0	650	0	0	1463	0	0	883	0
	$E(waste_t)$	0	0	518	0	0	500	0	0	563	0	0	283
Sim	α -serv.	100	99.0	95.2	100.0	98.6	95.2	100.0	100.0	95.1	100.0	100.0	95.0
	AvQ	2469	59	78	2140	205	36	2547	334	64	1132	157	64
	$\# orders$	1.00	0.05	0.05	0.92	0.15	0.04	0.86	0.14	0.12	0.78	0.15	0.10
	$Inv1$	1668	58	74	1331	164	34	2015	229	61	906	139	19
SDP	$Inv2$	0	720	57	26	599	150	11	1330	99	27	789	51
	$Waste$	0	0	525	14	1	465	54	0	564	57	11	323
	α -serv.	100	99.4	100.0	99.9	99.5	100.0	99.8	100.0	99.5	100.0	99.9	98.9

* MILP $E(TC) = 39192$; Simulated MILP Average $TC = 39231$; SDP Average $TC = 41269$.

2.6.1 Design of experiments and performance measures

The base case has a shelf life of $M = 3$ and a demand pattern that is erratic due to promotions in weeks 1, 2, 4, 5, 8 and 9, as depicted in Fig. 2.1. Table 2.7 reports the design of 86 experiments varying the parameters of the base case. Systematically we vary fixed setup cost k (1500, 500 and 2000), cost of waste w (-0.5, 0 and 0.5), α -service levels (90%, 95% and 98%), CV (0.1, 0.25 and 0.333) and shelf life M (2, 3, 4). The underlined values

are our base parameter values. The other cost values are constant: unit production cost $c = 2$ and unit holding cost $h = 0.5$. Note, negative cost of waste means the product has a salvage value, which is usually much less than the unit production cost c , zero cost of waste means that only the unit production cost are lost in case of waste, and positive cost of waste means that there is a cost to discard the wasted items.

Table 2.7 Design of Experiments

Experiment	Demand	k	w	α -service (%)	CV	M
Base	Erratic	1500	0	95	0.25	3
1 – 9	Erratic	1500	-0.5, 0, 0.5	90, 95, 98	0.10	3
*10 – 18	Erratic	1500	-0.5, 0, 0.5	90, 95, 98	0.25	3
19 – 27	Erratic	1500	-0.5, 0, 0.5	90, 95, 98	0.33	3
28 – 36	Erratic	500	-0.5, 0, 0.5	90, 95, 98	0.10	3
37 – 45	Erratic	500	-0.5, 0, 0.5	90, 95, 98	0.25	3
46 – 54	Erratic	500	-0.5, 0, 0.5	90, 95, 98	0.33	3
55 – 63	Erratic	2000	-0.5, 0, 0.5	90, 95, 98	0.10	3
64 – 72	Erratic	2000	-0.5, 0, 0.5	90, 95, 98	0.25	3
73 – 81	Erratic	2000	-0.5, 0, 0.5	90, 95, 98	0.33	3
82 – 83	Erratic	1500	0	95	0.25	2, 4
84	Err Variant	1500	0	95	0.25	3
85	Highly Err	1500	0	95	0.25	3
86	Stationary	1500	0	95	0.25	3

* including the base case

For the base parameter values we also tested three other demand patterns, depicted in Fig. 2.1. A variant of the erratic demand pattern is used, with different mean demands per period but the same overall mean and standard deviation (experiment 84). In this erratic variant a clustering of promotions can be observed in weeks 3 and 10. In experiment 85 a highly erratic demand pattern is tested, and in experiment 86 a stationary demand pattern. The total expected demand is 7200 for all patterns.

For the experiments listed in Table 2.7 the following performance indicators are measured by simulating the inventory system using the same (pseudo) random number series of 10,000 runs: expected total costs, total expected production quantity, expected waste as a percentage of the production quantity and the order policy denoting the periods with a production run. Moreover, an indicator Sum of Squared Errors of the α -service level $SSE(\alpha)$ measures the quality with respect to meeting the service level requirements: $SSE(\alpha) = \sum_t (\max\{0, \alpha - \text{realised service level in period } t\})^2$. A service level above the requirements is considered as good, so this is not considered an error. The lower the $SSE(\alpha)$, the better the service levels are met. This indicator is of main concern measuring how well the MILP model provides approximate solutions in terms of the service level. Each simulation run starts with

no inventory in stock and lasts $T = 12$ periods. For every period t , the reached α -service level is computed by calculating the fraction of runs in which the inventory at the end of period t is nonnegative.




Fig. 2.1 Demand patterns

2.6.2 Performance of the MILP approximation

First of all, we discuss the MILP production plan for the base case in detail and compare its results with a Stochastic Dynamic Programming (SDP) benchmark from Hendrix et al. (2012) designed for the same SP model. The SDP policy is not a replenishment cycle policy with fixed order periods Y_t but an order is placed given the state of the inventory levels of all ages. We call this $Q_t(I)$, where the argument of this function is the inventory at the end of last period that is not wasted, i.e. $I = (I_{1,t-1}, I_{2,t-1}, \dots, I_{M-1,t-1})$. This policy has more freedom to decide every period to order or not and how much, depending on the actual inventory levels, than the policy generated by the MILP model, where the timing and order-up-to levels are fixed. Therefore, we would expect the solution of SDP to be cheaper than the solution of MILP. SDP generates a policy that fulfils $P(I_{1t} \geq 0 \mid I) \geq \alpha$ for all $I = (I_{1,t-1}, I_{2,t-1}, \dots, I_{M-1,t-1})$.

This is a sufficient, but not necessary condition for chance constraint (4). SDP does not and cannot take the difference on probability of occurrence of all inventory states into account. That is the reason that the SDP overachieves the service levels at higher total costs, as noticed in Hendrix et al. (2012). The SDP generates a feasible, not necessarily optimal solution to the problem we defined. Table 2.8 shows the MILP production plan, the corresponding simulation results and the SDP benchmark. The base case is simulated for the MILP production plan and the SDP policy with 10,000 runs of the same random numbers.

Table 2.8 Base case: MILP production plan and MILP simulation results versus SDP results*

	T	1	2	3	4	5	6	7	8	9	10	11	12
Data	$E(d_t)$	800	950	200	900	800	150	650	800	900	300	150	600
MILP	S_t	1129	1550	600	2350	1450	650	1874	1224	1271	1333	1033	883
	$E(Q_t)$	1129	1221	0	1950	0	0	1874	0	847	962	0	0
	$E(I_{1t})$	329	600	0	1450	0	0	1224	0	371	962	0	0
	$E(I_{2t})$	0	0	400	0	650	0	0	424	0	71	883	0
	$E(waste_t)$	0	0	0	0	0	500	0	0	0	0	0	283
	S_t	1129	1550	600	2350	1450	650	1874	1224	1271	1333	1033	883
Sim. MILP	$Av(Q_t)$	1129	1221	0	1950	0	0	1880	0	848	975	0	0
	$Av(I_{1t})$	329	598	-5	1442	-2	-6	1225	-5	358	910	0	-11
	$Av(I_{2t})$	0	2	405	0	645	0	0	429	0	122	830	0
	$Av(waste_t)$	0	0	0	8	0	500	0	0	13	0	52	242
	α -serv.	95.0	99.5	95.3	100.0	98.6	95.1	100.0	95.3	95.0	100.0	100.0	89.0
	AvQ	2293	49	178	1872	153	91	1783	207	729	365	215	390
SDP	# orders	1.00	0.22	0.13	0.88	0.35	0.09	0.92	0.13	0.90	0.32	0.23	0.52
	$Inv1$	1493	45	170	1122	106	86	1271	169	438	341	199	118
	$Inv2$	0	548	37	57	425	81	29	535	66	195	304	64
	$Waste$	0	0	365	0	3	306	1	0	30	32	98	112
	α -serv.	100	98.8	99.9	99.8	98.9	99.9	99.9	99.5	99.0	99.7	99.9	96.6

* MILP $E(TC) = 28648$; Simulated MILP Average $TC = 28654$; SDP Average $TC = 29143$.

For the base case, the SDP policy overachieves the α -service level requirement in each period at a higher total cost of 29143. There is production in 5.7 periods on average, with a total average production quantity of 8327. The MILP order policy prescribes to order in 6 periods, namely 1, 2, 4, 7, 9 and 10 with expected total costs of 28648 and total expected order quantity of 7983. Simulating the MILP policy gives expected total costs of 28654 and a total expected order quantity of 8003. The MILP policy meets the α -service level of 95% in 11 periods. In period 12, the service level is only 89.0%, this may be well accepted in practice as production plans are usually updated every week.

Table 2.9 shows the results of the MILP policy of all 86 generated experiments. In the experiments, the setup cost of $k = 500$ gives, for a CV of 0.25 and 0.33, service levels close to the requirements. There are production runs in 9 of the 12 periods: only when demand is low, a production run is skipped. For a CV of 0.10, there are fewer production runs, sometimes resulting in no production in period 12, where the demand is high. The service level in period 12 is then only 78.3% or 80.4%. A setup cost of $k = 2000$ gives in most instances an order policy of ordering in periods 1, 4, 7 and 10. Considering the maximum shelf life, this is the minimum amount of production runs that is needed. That means that the inventory level before production starts is 0. There is no waste during the replenishment cycle due to inventory on hand. The service level requirements are closely met. In the experiments, the setup cost of $k = 1500$ gives mixed results with respect to meeting the

service level requirements, they are reasonable, with some ups and downs. The order policies prescribe more than the minimum amount of 4 production runs and less than 9, which results in a considerable amount of older items in stock.

Table 2.9 Results of the MILP model with a total demand of 7200 over 12 periods. The base case is marked grey.

Erratic demand pattern																		
CV = 0.10						CV = 0.25						CV = 0.33						
k=1500	E(TC)	Total E(Q)	% Waste	Order Policy	SSE(α)	E(TC)	Total E(Q)	% Waste	Order Policy	SSE(α)	E(TC)	Total E(Q)	% Waste	Order Policy	SSE(α)			
Service level 90%																		
w=-0.5	1	25057.5	7783	7.5	1.4.7.10	0.0	10	27210.5	8214	12.3	1.4.7.9.10	20.7	19	28172.5	8538	15.7	1.4.7.9.10	66.3
w=0	2	25349	7783	7.5	1.4.7.10	0.0	11	27717.5	8214	12.3	1.4.7.9.10	20.7	20	28748	8005	10.1	1.2.4.7.9.10	66.3
w=0.5	3	25583	7598	4.2	1.4.7.9.11	137.4	12	28176	7810	7.8	1.2.4.7.9.10	20.7	21	28912	7491	3.9	1.2.4.5.7.9.10	72.0
Service level 95%																		
w=-0.5	4	25467.5	7947	9.4	1.4.7.10	0.0	13	28062	8501	15.3	1.4.7.9.10	36.6	22	29335.5	8272	13.0	1.2.4.7.9.10	43.4
w=0	5	25841	7947	9.4	1.4.7.10	0.0	14	28648	7983	9.8	1.2.4.7.9.10	36.6	23	29606	7613	5.4	1.2.4.5.7.9.10	53.0
w=0.5	6	26050	7716	5.4	1.4.7.9.11	213.7	15	28835	7483	3.8	1.2.4.5.7.9.10	38.0	24	29812.5	7613	5.4	1.2.4.5.7.9.10	53.0
Service level 98%																		
w=-0.5	7	25932.5	8133	11.5	1.4.7.10	0.0	16	29045	8836	18.5	1.4.7.9.10	36.8	25	30429.5	7827	8.0	1.2.4.5.7.9.10	9.7
w=0	8	26383	7882	7.0	1.4.7.9.11	122.8	17	29357	7566	4.8	1.2.4.5.7.9.10	40.3	26	30743	7827	8.0	1.2.4.5.7.9.10	9.7
w=0.5	9	26660	7882	7.0	1.4.7.9.11	122.8	18	29540	7566	4.8	1.2.4.5.7.9.10	40.3	27	31056.5	7827	8.0	1.2.4.5.7.9.10	9.7
k=500	E(TC)	Total E(Q)	% Waste	Order Policy	SSE(α)	E(TC)	Total E(Q)	% Waste	Order Policy	SSE(α)	E(TC)	Total E(Q)	% Waste	Order Policy	SSE(α)			
Service level 90%																		
w=-0.5	28	19750	7401	1.7	1.2.4.5.7.9.12	0.0	37	20962	7393	0.0	1.2.4.5.7.9.10.12	0.0	46	21540.5	7454	0.0	1.2.4.5.7.8.9.10.12	0.0
w=0	29	19759.5	7280	0.0	1.2.4.5.7.9.11	137.4	38	20962	7393	0.0	1.2.4.5.7.9.10.12	0.0	47	21540.5	7454	0.0	1.2.4.5.7.8.9.10.12	0.0
w=0.5	30	19759.5	7280	0.0	1.2.4.5.7.9.11	137.4	39	20962	7393	0.0	1.2.4.5.7.9.10.12	0.0	48	21540.5	7454	0.0	1.2.4.5.7.8.9.10.12	0.0
Service level 95%																		
w=-0.5	31	20003	7309	0.1	1.2.4.5.7.9.11	213.8	40	21474.5	7447	0.0	1.2.4.5.7.8.9.10.12	0.1	49	22257.5	7565	0.5	1.2.4.5.7.8.9.10.12	0.1
w=0	32	20006.5	7309	0.1	1.2.4.5.7.9.11	213.8	41	21474.5	7447	0.0	1.2.4.5.7.8.9.10.12	0.1	50	22277	7565	0.5	1.2.4.5.7.8.9.10.12	0.1
w=0.5	33	20010	7309	0.1	1.2.4.5.7.9.11	213.8	42	21474.5	7447	0.0	1.2.4.5.7.8.9.10.12	0.1	51	22296.5	7565	0.5	1.2.4.5.7.8.9.10.12	0.1
Service level 98%																		
w=-0.5	34	20276.5	7324	0.0	1.2.4.5.7.9.10.12	0.1	43	22066.5	7522	0.2	1.2.4.5.7.8.9.10.12	0.1	52	23142	7768	2.1	1.2.4.5.7.8.9.10.12	0.1
w=0	35	20276.5	7324	0.0	1.2.4.5.7.9.10.12	0.1	44	22073	7522	0.2	1.2.4.5.7.8.9.10.12	0.1	53	23222.5	7768	2.1	1.2.4.5.7.8.9.10.12	0.1
w=0.5	36	20276.5	7324	0.0	1.2.4.5.7.9.10.12	0.1	45	22079.5	7522	0.2	1.2.4.5.7.8.9.10.12	0.1	54	23303	7768	2.1	1.2.4.5.7.8.9.10.12	0.1
k=2000	E(TC)	Total E(Q)	% Waste	Order Policy	SSE(α)	E(TC)	Total E(Q)	% Waste	Order Policy	SSE(α)	E(TC)	Total E(Q)	% Waste	Order Policy	SSE(α)			
Service level 90%																		
w=-0.5	55	27057.5	7783	7.5	1.4.7.10	0.0	64	29232.5	8653	16.8	1.4.7.10	0.0	73	30392.5	9117	21.0	1.4.7.10	0.0
w=0	56	27349	7783	7.5	1.4.7.10	0.0	65	29959	8653	16.8	1.4.7.10	0.0	74	31341.5	8538	15.7	1.4.7.9.10	66.3
w=0.5	57	27640.5	7783	7.5	1.4.7.10	0.0	66	30685.5	8653	16.8	1.4.7.10	0.0	75	32010.5	8538	15.7	1.4.7.9.10	66.3
Service level 95%																		
w=-0.5	58	27467.5	7947	9.4	1.4.7.10	0.0	67	30260	9064	20.6	1.4.7.10	0.0	76	31747.5	9659	25.5	1.4.7.10	0.0
w=0	59	27841	7947	9.4	1.4.7.10	0.0	68	31192	9064	20.6	1.4.7.10	0.0	77	32750	8955	19.6	1.4.7.9.10	43.3
w=0.5	60	28214.5	7947	9.4	1.4.7.10	0.0	69	31863	8501	15.3	1.4.7.9.10	36.6	78	33312.5	7613	5.4	1.2.4.5.7.9.10	53.0
Service level 98%																		
w=-0.5	61	27932.5	8133	11.5	1.4.7.10	0.0	70	31415	9526	24.4	1.4.7.10	0.0	79	33275	10270	29.9	1.4.7.10	0.0
w=0	62	28399	8133	11.5	1.4.7.10	0.0	71	32363	8836	18.5	1.4.7.9.10	36.8	80	34243	7827	8.0	1.2.4.5.7.9.10	9.7
w=0.5	63	28865.5	8133	11.5	1.4.7.10	0.0	72	33040	7566	4.8	1.2.4.5.7.9.10	40.3	81	34556.5	7827	8.0	1.2.4.5.7.9.10	9.7
k=1500 w=0 Service level 95%																		
Total %																		
E(TC) E(Q) Waste Order Policy SSE(α)																		
M=2	30787.5 9005 20.0 1.2.4.5.7.9.11 0.1																	
M=4	26807 7666 6.1 1.4.7.9 9.1																	
EV	27719.5 8093 11.0 1.3.5.7.10 345.2																	
HE	27400 8163 11.8 1.4.6.7.10 529.2																	
ST	27992 7549 0.0 1.3.5.7.9.11 1.5																	

is overestimated. Second, the combination of demand pattern and parameter values may result in inconvenient production moments. A replenishment cycle can be of different length, varying from 1 period to M periods. If the current replenishment cycle is 1 period, there can only be waste at the end of the period. This will not affect the next period, because a new replenishment cycle will start, and the waste will not be part of the inventory on hand. If the current replenishment cycle will be more than one period, during the replenishment cycle there can be waste. If the amount of waste during the replenishment cycle is more than expected, the service level might not be met. Over all the performed experiments, in 3.6% of the periods the realised service level is more than 1% lower than the required service level. This happened only in period 12. The MILP approximation provides a practical solution to the SP problem. As production plans are updated frequently, not meeting the service level in the last period(s) is less relevant.

The performance of the MILP solutions with respect to the required service level becomes more complicated when there are many items (of different ages) in stock. The simulation determines the order quantity as the order-up-to level of the MILP model minus the on-hand inventory level of ages one and two. The results show that in the application of the model, it is important to investigate the need of considering the age-distribution of the items in stock for the determination of the order quantity.

In summary, the results of the simulation of the MILP policy show that the quality of the solution is data-dependent. There are three situations with data such that the production plan prescribes:

1. Order every period: the service levels are fulfilled;
2. Order every M periods: the service levels are fulfilled;
3. Order irregular: the service levels are fulfilled with an error tolerance of 1%, except for end-of-horizon effects.

In situations 1 and 2 the MILP approximation is fast with very good results. In situation 3, the MILP model gives a lower bound solution but needs correction to meet all service levels with an even smaller error tolerance. In the experiments the focus is on case 3.

Table 2.10 shows a summary of the effect of varying the parameters on the solutions of the MILP model. The effect of varying the α -service level or the CV on the expected total costs follows the intuition: the higher the service level or the CV, the higher the expected total costs $E(TC)$. However, in the experiments, the effect on the total expected order quantity $E(Q)$ and the percentage of waste is not unidirectional, because the number of production runs as well as the timing of production runs may change. More production runs lead to a lower total expected order quantity and less waste. A higher cost of waste or a lower setup cost may increase the number of production runs. A longer shelf life reduces the number of production runs and therefore the expected total costs, order quantity and waste, in line with the intuition.

Table 2.10 Summary of the parameter analysis

Vary:	Observed effect on:			
	$E(TC)$	Total $E(Q)$	% Waste	# production runs
α -service level \uparrow	\uparrow	\uparrow or \downarrow	\uparrow or \downarrow	\uparrow or $=$
CV \uparrow	\uparrow	\uparrow or \downarrow	\uparrow or \downarrow	\uparrow or $=$
Cost of waste \uparrow	\uparrow or $=$	\downarrow or $=$	\downarrow or $=$	\uparrow or $=$
Setup cost \downarrow	\downarrow	\downarrow or $=$	\downarrow or $=$	\uparrow or $=$
Shelf life \uparrow	\downarrow	\downarrow	\downarrow	\downarrow

Fig. 2.2 shows the influence of the α -service level and the cost of waste on the number of production runs en percentage of waste. This illustrates that the model can be used to manage the amount of waste, while maintaining a certain service level.




Fig. 2.2 The case of erratic demand and CV = 0.25: influence of service level on number of production runs and percentage of waste; influence of the cost of waste on percentage of waste at different service levels. The numbers on top of the bars denote the number of production runs

2.6.3 Benchmark with SDP

For the experiments with maximum shelf life $M = 3$, a comparison is made between the simulated MILP solutions and the SDP solutions. In Table 2.11 a summary of these results is presented. On average the SDP production plan is 1.2% more costly than the MILP production plan. In the SDP solutions the average total order quantity is 0.8% less and the percentage of waste is 16.7% lower than in the MILP solutions. The number of production runs is 9.3% higher. Because SDP policies tend to produce more often, SDP performs best, with respect to the cost, when the setup cost k is low. The $SSE(\alpha)$ of the SDP solutions is always 0 or not significantly above 0, therefore it is not presented. Appendix 2.B provides a table summarising the results of the MILP policy compared to that of the SDP policy for all experiments.

Table 2.11 Percentual Average Total Costs, Order Quantity and Waste of the SDP policy in comparison with the simulation results of the MILP policy

Parameter	Average Total Costs 100*SDP/MILP	Average Total Order Quantity 100*SDP/MILP	Average Total Waste 100*SDP/MILP	Average # Production runs 100*SDP/MILP	Avg # Pr.runs MILP	Avg # Pr.runs SDP
TotalAvg	101.2	99.2	83.3	109.3	6.2	6.8
$\alpha = 90$	102.4	99.2	85.1	114.9	5.8	6.7
$\alpha = 95$	101.2	99.2	85.2	109.5	6.2	6.8
$\alpha = 98$	100.0	99.3	79.1*	104.1	6.6	6.9
CV = 0.1	101.4	99.2	90.0*	113.8	5.3	6.0
CV = 0.25	101.3	99.0	79.9	110.1	6.4	7.0
CV = 0.33	100.9	99.5	80.7	105.0	7.0	7.3
$w = -0.5$	101.5	98.2	85.2*	112.7	5.8	6.5
$w = 0$	101.1	99.5	93.0*	108.6	6.2	6.8
$w = 0.5$	100.9	100.0	92.5*	106.9	6.6	7.0
$k = 500$	99.8	99.8	92.2*	106.7	8.3	8.9
$k = 1500$	101.4	99.7	76.0	108.3	5.6	6.1
$k = 2000$	102.3	98.2	82.6	115.1	4.6	5.3

* In experiments 34, 35 and 36 the MILP solutions have an average total waste of 0.1, resulting in high numbers in comparison with SDP. These values are excluded from the average total waste percentages.

2.7 Conclusions

We studied the practical production planning problem of a food producer facing a non-stationary erratic demand for a perishable product with a fixed life time, under a service-level constraint. The case includes a fixed setup cost for every production run, zero lead time and a First In First Out issuing policy. In case of out-of-stock, demand is backlogged. A theoretical Stochastic Programming model for this problem has been presented, that considers the age-distribution of the items in stock. The question is how to generate a waste-compensating replenishment cycle policy by applying commercial MILP solvers. Therefore, an MILP model has been formulated to generate approximate solutions. A solution provides a plan specifying simultaneously the periods to produce and the corresponding order-up-to-levels. To meet a certain α -service level, the model considers and corrects for the expected age-distribution of the items in stock. The model can be solved in a fraction of a second. This makes it interesting for practical applications.

In situations where the replenishment cycles are always one period, or always equal to the maximum shelf life, the service levels are fulfilled. The MILP model is specifically tested on

the production plans with replenishment cycles of irregular lengths. In this situation, the MILP solutions provide approximations of the required service level. Simulation shows that in 96.4% of the periods the service level requirements are met, with an error tolerance of 1%. The performance of the MILP model with respect to service level requires attention when there are many items of different ages in stock. In the application of the model, it is important to investigate the need of considering the age-distribution of the items in stock for the determination of the order quantity.

Appendix 2.A

Table 2.A Used symbols, where applicable: $E(\cdot)$ = expected value of...

SP	MILP	
T	T	maximum number of periods of the finite time horizon
t	t	index denoting the period
M	M	fixed maximum (internal) shelf life
b	b	index denoting the age of the item
k	k	fixed setup cost for every production run
c	c	variable production cost per item produced
h	h	holding cost per item, for items that are carried over from one period to the next
w	w	disposal cost per item of waste
α	α	service level
	M	big number
d_t	$E(d_t)$	non-stationary stochastic demand
$g_t(\cdot)$	$g_t(\cdot)$	probability density function of demand d_t
$G_t(\cdot)$	$G_t(\cdot)$	cumulative distribution function of demand d_t
I_{bt}	$E(I_{bt})$	the inventory level of items with age b at the end of period t
I_{Mt}	$E(I_{Mt})$	inventory of age M at the end of period t is considered waste
Q_t	$E(Q_t)$	production quantity in period t
S_t	S_t	the order-up-to level or target starting inventory level at the beginning of period t
$E(TC)$	$E(TC)$	expected total costs over the time horizon
Y_t	Y_t	binary variable takes the value of 1 if there is a production run in period t , and 0 otherwise
	Z_{tj}	binary variable takes the value of 1 if the most recent order prior to period t was in period $t - j + 1$, and 0 otherwise
	$E(X_{bt})$	auxiliary variable denotes the residual demand for items of age b with $b = 1, \dots, M - 1$ in period t
	BX_{bt}	binary variable takes the value of 1 if $E(X_{bt}) > 0$, and 0 otherwise

Appendix 2.B

Table 2.B Comparison between the simulated MILP solutions and the SDP solutions

Exp.	100*SDP/MILP Total Costs	100*SDP/MILP Total Q	MILP % waste	SDP % waste	MILP # Pr. runs	SDP # Pr. runs	MILP SSE(α)*
1	102.5	98.3	7.8	5.2	4	5.1	0.0
2	102.1	98.1	7.8	4.9	4	5.1	0.0
3	101.9	99.8	4.7	4.1	5	5.3	137.4
4	101.6	97.2	9.5	5.7	4	5.1	0.0
5	100.9	97.0	9.5	5.5	4	5.1	0.0
6	100.9	98.9	5.9	4.5	5	5.3	213.7
7	100.6	95.9	11.5	6.2	4	5.1	0.0
8	99.8	98.1	7.3	5.3	5	5.3	122.8
9	99.5	97.7	7.3	5.0	5	5.3	122.8
10	103.2	100.3	12.9	11.2	5	5.5	20.7
11	102.7	99.2	12.9	10.2	5	5.6	20.7
12	102.4	99.7	8.3	6.0	6	6.5	20.7
13	102.0	98.8	15.7	12.4	5	5.5	36.6
14	101.7	104.0	10.2	11.4	6	5.7	36.6
15	101.6	102.1	4.4	3.9	7	7.1	38.0
16	100.7	97.7	18.9	14.3	5	5.5	36.8
17	100.9	102.1	5.8	4.7	7	7.1	40.3
18	100.7	101.6	5.8	4.3	7	7.2	40.3
19	103.7	99.5	16.7	13.9	5	5.7	66.3
20	103.1	98.0	11.1	6.7	6	6.9	66.3
21	102.8	102.4	5.5	5.2	7	7.2	72.0
22	102.0	98.8	14.0	10.0	6	6.6	43.4
23	101.6	101.9	7.6	6.4	7	7.2	53.0
24	101.3	101.3	7.6	6.0	7	7.3	53.0
25	100.0	103.6	11.0	10.5	7	6.8	9.7
26	99.5	99.9	11.0	7.2	7	7.3	9.7
27	99.0	99.5	11.0	6.9	7	7.3	9.7
28	101.0	99.2	1.7	0.8	7	8.2	0.0
29	101.2	100.8	0.4	0.7	7	8.2	137.4
30	101.3	100.6	0.4	0.5	7	8.3	137.4
31	100.7	100.9	0.6	0.9	7	8.2	213.8
32	100.8	100.7	0.6	0.7	7	8.3	213.8
33	100.8	100.6	0.6	0.6	7	8.3	213.8
34	99.9	101.0	0.0	1.0	8	8.2	0.1
35	100.1	100.9	0.0	0.8	8	8.3	0.1
36	100.3	100.7	0.0	0.6	8	8.3	0.1
37	100.5	99.8	0.7	0.5	8	9.5	0.0
38	100.5	99.8	0.7	0.5	8	9.6	0.0

Chapter 2

Exp.	100*SDP/MILP Total Costs	100*SDP/MILP Total Q	MILP % waste	SDP % waste	MILP # Pr. runs	SDP # Pr. runs	MILP SSE(α)*
39	100.5	99.8	0.7	0.5	8	9.6	0.0
40	99.7	99.7	1.1	0.9	9	9.4	0.1
41	99.7	99.7	1.1	0.9	9	9.4	0.1
42	99.6	99.7	1.1	0.9	9	9.4	0.1
43	98.9	99.5	1.9	1.6	9	9.1	0.1
44	98.9	99.5	1.9	1.6	9	9.1	0.1
45	98.8	99.5	1.9	1.6	9	9.1	0.1
46	100.0	99.5	2.2	1.7	9	9.4	0.0
47	99.9	99.5	2.2	1.7	9	9.4	0.0
48	99.8	99.5	2.2	1.7	9	9.4	0.0
49	99.2	99.2	3.4	2.7	9	9.1	0.1
50	99.1	99.2	3.4	2.7	9	9.1	0.1
51	99.0	99.2	3.4	2.7	9	9.1	0.1
52	98.4	98.9	5.1	4.1	9	8.8	0.1
53	98.3	98.9	5.1	4.1	9	8.8	0.1
54	98.1	98.9	5.1	4.1	9	8.8	0.1
55	103.3	101.4	7.8	8.8	4	4.3	0.0
56	103.5	100.5	7.8	7.9	4	4.4	0.0
57	103.4	99.9	7.8	7.3	4	4.5	0.0
58	102.4	100.3	9.5	9.6	4	4.3	0.0
59	102.4	99.6	9.5	8.8	4	4.4	0.0
60	102.2	97.0	9.5	5.5	4	5.1	0.0
61	101.6	99.3	11.5	10.6	4	4.3	0.0
62	101.5	98.7	11.5	10.0	4	4.4	0.0
63	100.8	95.5	11.5	5.8	4	5.1	0.0
64	105.0	97.1	17.3	12.8	4	5.2	0.0
65	104.1	95.8	17.3	11.8	4	5.4	0.0
66	103.1	94.9	17.3	11.0	4	5.5	0.0
67	103.4	94.4	20.8	13.8	4	5.3	0.0
68	102.1	93.1	20.8	12.7	4	5.5	0.0
69	101.7	98.4	15.7	11.9	5	5.6	36.6
70	101.8	92.1	24.5	15.4	4	5.4	0.0
71	100.8	98.2	18.9	14.7	5	5.5	36.8
72	100.4	106.7	5.8	8.8	7	6.4	40.3
73	105.1	95.7	21.7	15.9	4	5.4	0.0
74	104.0	100.1	16.7	14.4	5	5.6	66.3
75	103.5	98.9	16.7	13.3	5	5.7	66.3
76	103.1	92.7	25.7	17.3	4	5.5	0.0
77	102.1	98.1	20.5	16.2	5	5.6	43.3
78	101.3	105.8	7.6	9.7	7	6.6	53.0
79	101.2	90.6	30.0	19.4	4	5.5	0.0

MILP approximation generating a YS policy

Exp.	100*SDP/MILP Total Costs	100*SDP/MILP Total Q	MILP % waste	SDP % waste	MILP # Pr. runs	SDP # Pr. runs	MILP SSE(α)*
80	99.5	105.4	11.0	11.9	7	6.5	9.7
81	99.4	101.0	11.0	8.0	7	7.1	9.7
Avg	101.2	99.2	8.8	6.8	6.2	6.8	
EV	102.2	96.5	12.3	8.4	5	5.8	345.2
HE	102.3	97.3	12.8	9.2	5	5.7	529.2
ST	104.5	101.9	1.3	4.2	6	6.2	1.5

* The SSE(α) of the SDP solutions is always 0 or not significantly above 0, therefore it is not presented.

Chapter 3

SDP in inventory control: a $Q(X)$ policy

Based on:

Pauls-Worm, K.G.J., Hendrix, E.M.T., 2015. SDP in Inventory Control: Non-stationary Demand and Service Level Constraints, in: Gervasi, O., Murgante, B., Misra, S., Gavrilova, M.L., Rocha, A.M.A.C., Torre, C., Tanar, D., Apduhan, B.O. (Eds.), Computational Science and Its Applications -- ICCSA 2015. Springer International Publishing, pp. 397-412.

Abstract

Inventory control implies dynamic decision making. Therefore, dynamic programming seems an appropriate approach to look for order policies. The implementation of service level requirement provides a big challenge. In this chapter, small instances are constructed to study the behaviour for an α -service level constraint and a fill rate constraint per period, being minimal service level constraints. To get a clear view on the system dynamics of the service level constraints, a non-perishable product is considered, to avoid the complication of the age-distribution of perishable items in stock. Evaluation of the SDP order policies shows that SDP generates order policies that reach a higher service level than required, because the service level requirement is dealt with in a conditional way depending on the current state of inventory. The SDP policy is compared with an order-up-to level policy. The order-up-to levels were generated by full enumeration meeting the service level requirement. This results in a policy with lower costs and service levels, which are at the end of the replenishment cycle close to the required service level. In general, one can conclude that SDP generates order policies that meet a conditional service level constraint. Given a minimal expected service level requirement, the resulting policy is not necessarily optimal.

3.1 Introduction

This study is motivated by practical cases of inventory management. In the practical case, a retailer or a producer faces non-stationary demand for a product and has to determine when and how much to order or to produce to meet a certain service level. The decision maker, hereafter manager, uses a periodic review, meaning that on fixed moments in time, e.g. every day or every week, the manager decides on the order quantity. Many products in retail face a non-stationary demand (Chapter 2), (Pauls-Worm et al., 2014). The decision on the order quantity is inherently a multistage problem, and therefore Stochastic Dynamic Programming (SDP) seems an appropriate approach to attempt to solve the problem (Bellman and Lee, 1978). An earlier application of SDP to a problem with non-stationary demand and a service level constraint (Hendrix et al., 2012) has shown that it does not necessarily generate an optimal solution for this case. This motivates a further study on this phenomenon.

SDP has been used for perishable products with stationary demand by Van Zyl (1964) to derive an optimal order-up-to policy. His work was extended by Fries (1975), who also used SDP. Sobel and Zhang (2001) studied a system for non-perishable products with combined variable deterministic and non-stationary stochastic demand using SDP to find an optimal policy. The above models typically use penalty costs when demand exceeds the stock level. In practice however, often a service level constraint is imposed. Chen and Krass (2001) show that in general it is not possible to transform a service level model into a cost model, above a certain “critical” service level.

In this chapter, we consider two types of service levels used in practice, the α -service level, also called cycle service level, and the β -service level, also called fill rate. Chopra and Meindl (2010) define cycle service level as the probability of not having a stock-out in a replenishment cycle. A replenishment cycle is the time between two orders. The fill rate indicates that a predefined percentage of the demand per replenishment cycle has to be fulfilled from stock. Chen and Krass (2001) define the difference between mean service level constraints and minimal service level constraints. Mean service level constraints measure the service level over the time horizon, whereas minimal service level constraints measure the service level in every period. Van Houtum and Zijm (2000) show for inventory systems with (non-)stationary demand the conditions for a one-to-one relation between cost and service models, where they consider a mean service level. Bijvank and Vis (2012) formulated a lost-sales inventory model with an average fill rate requirement as a constrained dynamic programming problem for stationary demand. They solved the problem with Lagrange relaxation by a value-iteration algorithm to find optimal replenishment policies. In this chapter we focus on minimal service level constraints for the non-stationary demand case and we investigate whether SDP is a suitable method if a service level constraint applies. One stylized example is elaborated for different situations. It shows how in general SDP handles service level constraints and does not necessarily provide the optimal solution.

Section 3.2 introduces a stylized example of variable demand. It has a finite time horizon of 6 periods representing the possibility of a manager in retail to order at most six times a week. This simple instance shows the application of Dynamic Programming (DP) to solve the

deterministic variable discrete demand case. In Section 3.3 the example is extended to a stochastic uniform discrete demand case and solved with SDP to generate $Q_t(I)$ order policies, where the order quantity Q depends on the inventory level I on hand. Section 3.4 adds an α -service level constraint and Section 3.5 a fill rate constraint. The derived SDP solution for both cases is compared to an optimal policy fulfilling the service level constraints found by full enumeration. Section 3.6 illustrates the SDP approach for a continuous distribution of demand by a Gamma distribution. Section 3.7 concludes Chapter 3.

3.2 Dynamic Programming

We consider a small inventory problem with a deterministic variable demand. The problem has a finite time horizon of $T = 6$ and demand is $d_t = 3, 1, 2, 4, 3, 2$. The setup cost is $k = 5$ and holding cost $h = 1$. Holding cost is paid over the inventory that is carried over to the next period. The unit variable cost is $c = 0$ and the starting inventory $I_0 = 0$. An order is placed and delivered at the beginning of the period, demand occurs during the period, and the inventory level is calculated at the end of the period. The question is when to order and how much. This problem can be formulated as follows:

$$\text{Min } TC = \sum_{t=1}^T g(Q_t) + hI_t \quad (1)$$

where procurement cost is given by the function

$$g(x) = k + cx, \text{ if } x > 0, \text{ and } g(0) = 0 \quad (2)$$

Subject to

$$I_t = I_{t-1} + Q_t - d_t \quad t = 1, \dots, T \quad (3)$$

$$I_0 = 0 \quad (4)$$

$$I_t, Q_t \geq 0 \quad t = 1, \dots, T \quad (5)$$

The objective function (1) represents the total relevant costs, consisting of ordering cost and holding cost. In (2) the ordering costs are specified by a fixed setup cost when an order is placed and variable procurement cost. Equations (3) are balance equations stating that the inventory level at the end of period t equals the inventory level at the end of the period before, increased with the order quantity minus the demand in period t where (4) gives the initial inventory level. Equations (5) are non-negativity constraints.

Wagner and Whitin (1958) derived properties of the optimal solution to this problem and developed an efficient algorithm to solve the problem. The algorithm they developed is a modification of Dynamic Programming. This paper applies DP in order to develop towards SDP with service level constraints.

In the example, the first order Q_1 fulfills the accumulated total demand for a number of future periods from one to six, i.e. $Q_1 = 3, 4, 6, 10, 13$ or 15 . After demand $d_1 = 3$ is realized, the possible inventory level is $I_1 = 0, 1, 3, 7, 10$ or 12 . In period $t = 2$, demand $d_2 = 1$ will be fulfilled from the available inventory, or, in case $I_1 = 0$, there will be an order to fulfill demand of period 2 up to 6, or to fulfill demand of period 2 up to 5, 4 or 3 or just

period 2. The same reasoning holds for periods 3 to 5, so either $Q_t = 0$, or $Q_t = \sum_{j=t}^N d_j$ for some N , $t \leq N \leq T$.

The beginning of a period, when the decision about the order quantity has to be made, is called a stage in DP. A stage starts with the inventory level at the end of the previous period which is called the state. The set of possible values is called the state space. DP has the property that the optimal decision on time t (stage t) with state I_{t-1} , does not depend on the decisions made leading to state I_{t-1} . This means that the problem can be decomposed into sub-problems. The stages and states of this small instance are depicted in Fig. 3.1. A path from the start (stage 1) to the end of period 6 models the timing and quantity of ordering and the inventory levels at the end and beginning of each period (stage). Given the possibility to order for all future periods, we have $T! = 720$ possible paths for this inventory problem. DP generally starts with a backward procedure, followed by a forward procedure.

The backward procedure assigns a valuation to each state, called a value function $V_t(I)$ to each state in the system. Stage 6 ends with $I_6 = 0$ at the end of the time horizon. $V_6(I)$ equals the costs that have to be made from stage t at state I to the end of the final stage 6 with state $I = 0$. The so-called recursive Bellman equation for the value function $V_t(I)$ can be written as

$$V_t(I) = \min_Q \{g(Q) + h \cdot (I + Q - d_t) + V_{t+1}(I + Q - d_t)\} \text{ for } t = T, \dots, 1 \quad (6)$$

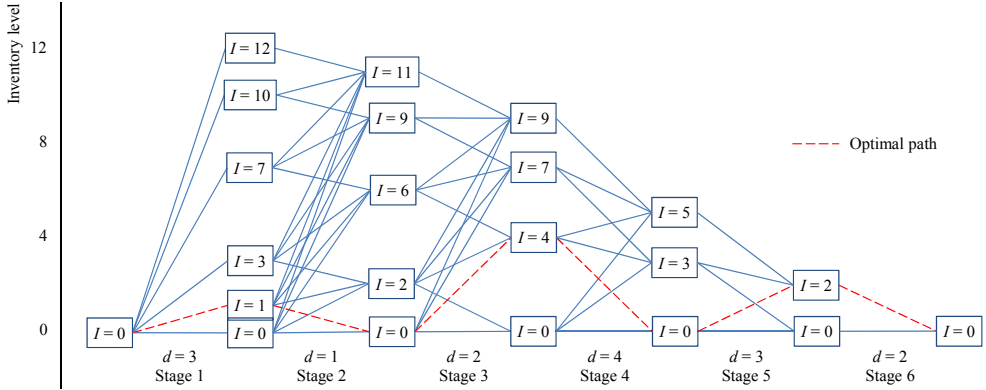


Fig. 3.1 All possible states of inventory levels and all possible paths

We illustrate the backward procedure for this small instance. Consider $V_6(2) = 0$. Holding cost is paid at the end of the previous period, so no cost has to be paid in this state; no order takes place. $V_6(0) = k + c \cdot Q_6 = 5 + 0 = 5$. An order of $Q_6 = d_6 = 2$ is placed at a cost of $k = 5$.

$V_5(5) = h \cdot (5 + 0 - 3) + V_6(5 + 0 - 3) = 2 + 0 = 2$. No order takes place, but at the end of stage 5, there is an inventory level of 2 that incurs holding cost.

$V_5(3) = \min_{Q \in \{0,2\}} \{h(3 + Q - 3) + V_6(3 + Q - 3)\} = \min\{(0+5), (7+0)\} = 5$. For $I_4 = 3$, the choice is not to order, or to order for demand of period 6. The minimum cost corresponds to not ordering.

$V_5(0) = \min_{Q \in \{3,5\}} \{k + h(0 + Q - 3) + V_6(0 + Q - 3)\} = \min\{5+0+5, 5+2+0\} = 7$. Here the minimum costs are obtained for $Q_5 = d_5 + d_6$.

Continuing this process provides finally the minimum total cost of $V_1(0) = 22$.

The path leading to the minimum cost can be determined in a forward procedure. The minimum cost at $V_1(0)$ is obtained via $I_1 = 1$, $I_2 = 0$, $I_3 = 4$, $I_4 = 0$, $I_5 = 2$ and $I_6 = 0$, so the optimal ordering path is $Q_1 = 4$, $Q_2 = 0$, $Q_3 = 6$, $Q_4 = 0$, $Q_5 = 5$ and $Q_6 = 0$ represented by the dashed red path in Fig. 3.1.

3.3 Stochastic Dynamic Programming

We continue the example from Section 3.2, but now demand is stochastic and non-stationary, with expected demand $\mu_t = 3, 1, 2, 4, 3, 2$. Demand has a discrete Uniform distribution with $d_t \sim U\{0, 2\mu_t\}$. All demand has to be met. Notice that one can set this target as the support of the distribution is finite. Demand in period $t = 1$ can take the values $d_{1j} \in \{0, 1, 2, 3, 4, 5, 6\}$, where the index is $j = 1, \dots, N_t$ with $N_t = (2\mu_t + 1)$. Every outcome d_{tj} has probability of occurrence $p_t = 1/N_t$ so $p_1 = 1/7$, $p_2 = 1/3$, etc. The maximum possible

inventory level equals $I_{\max} = \sum_{t=1}^T 2\mu_t = 30$. In this model, this number bounds the state

space.

In the optimal solution of the deterministic example, ordering only takes place when the inventory level equals zero, and the inventory level is always the sum of demand in upcoming periods. In this stochastic example demand is uncertain, so ordering might be necessary even if the inventory is not zero. That results in many more possible inventory paths. The value function for this problem is defined by

$$V_t(I) = \min_{Q \in F_t(I)} \left[g(Q) + h \cdot p_t \sum_{j=1}^{N_t} (I + Q - d_{tj})^+ + p_t \sum_{j=1}^{N_t} V_{t+1}(I + Q - d_{tj})^+ \right] \quad (7)$$

$$F_t(I) = [0, (d_{tN_t} + \dots + d_{tN_T} - I)^+] \quad (8)$$

where $F_t(I)$ is the feasible area of Q and $x^+ = \max\{x, 0\}$. The optimal order quantity as function of inventory level I is

$$Q_t(I) = \arg \min_{Q \in F_t(I)} \left[g(Q) + h \cdot p_t \sum_{j=1}^{N_t} (I + Q - d_{tj})^+ + p_t \sum_{j=1}^{N_t} V_{t+1}(I + Q - d_{tj})^+ \right] \quad (9)$$


The SDP approach has been implemented in Matlab. Table 3.1 shows the generated optimal order quantity $Q_t(I)$ for each stage and inventory.

One can observe that the optimal order quantities follow in fact from order-up-to levels, i.e. there are values S_t such that $Q_t(I) = (S_t - I)^+$. As all demand has to be fulfilled, the order-up-to levels have the size $S_t = 2\mu_t$. The expected total costs are $V_1(0) = 38.49$. The introduction of uncertainty gives a 75% cost increase when compared to the corresponding

Table 3.1 SDP order quantities given the inventory level

I_{t-1}	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6
0	6	2	4	8	6	4
1	0	1	3	7	5	3
2	0	0	2	6	4	2
3	0	0	1	5	3	1
4	0	0	0	4	2	0
5	0	0	0	3	1	0
6	0	0	0	2	0	0
7	0	0	0	1	0	0
8	0	0	0	0	0	0

deterministic variant of the problem. The objective function value of expected costs of the optimal order quantities $Q_t(I)$ follows from the generation of all possible demand paths as illustrated in Fig. 3.2. There are $N_1 \cdot N_2 \cdot N_3 \cdot N_4 \cdot N_5 \cdot N_6 = 33075$ possible paths. The evaluation computes exactly the same expected total costs of $E(TC) = 38.49$ as the value of $V_1(0)$ in the SDP approach.

**Fig. 3.2** Possible inventory development given demand paths

3.4 SDP with an α -service level constraint

In Section 3.3, 100% of the demand has to be fulfilled. Now we add an α -service level constraint to the problem. We define the α -service level as a minimum probability α of not being out of stock at the end of a period, so

$$P(I_t \geq 0) \geq \alpha \quad t = 1, \dots, T \quad (10)$$

This means that at the start of a period holds

$$P(I_{t-1} + Q_t - d_t \geq 0) \geq \alpha \quad t = 1, \dots, T$$

$$\text{or } P(d_t \leq I_{t-1} + Q_t) \geq \alpha \quad t = 1, \dots, T \quad (11)$$

The total quantity $I_{t-1} + Q_t$ can be seen as a basic order quantity, which we will call \hat{Q}_t . The cumulative distribution function of $d \sim U\{a, b\}$ is $F(k) := P(d \leq k) = \frac{\lfloor k \rfloor - a + 1}{b - a + 1}$ for any $k \in [a, b]$. For $d_t \sim U\{a, b\}$ we have $a = 0$, $b = 2\mu_t$ and $b - a + 1 = N_t = (2\mu_t + 1)$. Elaboration of (11) for this case gives

$$P(d_t \leq \hat{Q}_t) = \frac{\lfloor \hat{Q}_t \rfloor + 1}{N_t} \geq \alpha \quad t = 1, \dots, T \quad (12)$$

$$\text{This results in } \hat{Q}_t = \lfloor \alpha \cdot N_t \rfloor \quad t = 1, \dots, T \quad (13)$$

The basic order quantity \hat{Q}_t is a lower bound for the order quantity, or the starting inventory in a period, necessary to meet the α -service level requirement.

Consider the case of Section 3 with $k = 5$ or 50 , $\alpha = 0.8$ and expected demand $\mu_t = 3, 1, 2, 4, 3, 2$. Recursion (7) is still valid, but the feasible area is now

$$F_t(I) = [\hat{Q}_t, (d_{tN_t} + \dots + d_{TN_t} - I)^+] \quad (14)$$

Demand that is not met will be lost. An SDP approach has been implemented in MATLAB. The results of the SDP approach and the evaluated service levels are listed in Table 3.2.

Table 3.2 SDP order policy and its reached service levels (required $\alpha = 0.8$)

$k = 5$; E(TC) = 36.95							$k = 50$; E(TC) = 129.01						
Service	0.86	1.00	1.00	0.89	0.89	1.00	Service	1.00	1.00	1.00	0.999	0.989	1.00
\hat{Q}_t	5	2	4	7	5	4	\hat{Q}_t	5	2	4	7	5	4
I_{t-1}	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6	I_{t-1}	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6
0	5	2	4	7	5	4	0	18	16	16	15	10	4
1	0	1	3	6	4	3	1	0	15	15	14	9	3
2	0	0	2	5	3	2	2	0	0	14	13	8	2
3	0	0	1	4	2	1	3	0	0	13	12	7	1
4	0	0	0	3	1	0	4	0	0	0	11	6	0
5	0	0	0	2	0	0	5	0	0	0	10	0	0
6	0	0	0	1	0	0	6	0	0	0	9	0	0
7	0	0	0	0	0	0	7	0	0	0	0	0	0
8	0	0	0	0	0	0	:	0	0	0	0	0	0
9	0	0	0	0	0	0	18	0	0	0	0	0	0

Table 3.2 shows that in case of $k = 5$ one orders on average (almost) every period where Q_t at inventory level $I_{t-1} = 0$ equals the basic order quantity \hat{Q}_t . Basic order quantity \hat{Q}_t and the order quantities Q_t behave as order-up-to levels. In case $k = 50$, one can observe in the different periods different reorder points. When the inventory level is at the reorder point or lower, one orders up to the level $Q_t(0)$, otherwise there is no order. E.g. in period 4, the reorder point equals 3. If $I = 4$ or more, this is enough to meet demand and service level in period 4. If $I = 3$ or less, there is an order to meet demand and service level in period 4 and the upcoming periods. In three of the six periods, the service levels are slightly higher than

required. The difference is due to the discrete and small numbers of demand. In the other periods, as well as in all periods of the case of $k = 50$, the evaluated service levels are much higher. Fig. 3.3 shows inventory development for all possible demand paths. Also inventory levels below zero (lost sales) are depicted. These are virtual inventory levels to show out-of-stocks, but lost sales is assumed, so the inventory level never drops below zero. Notice that the different inventory levels have different chances of occurrence. From all possible starting inventory levels, no matter how small the chance of occurrence, SDP wants to meet the α -service level requirement. This results in an overall overachievement of the service level. The SDP approach due to the definition of the feasible area tightens the α -service level constraint in (10) towards the conditional variant

$$P(I_t \geq 0 | I_{t-1}) \geq \alpha \quad t = 1, \dots, T \quad (15)$$

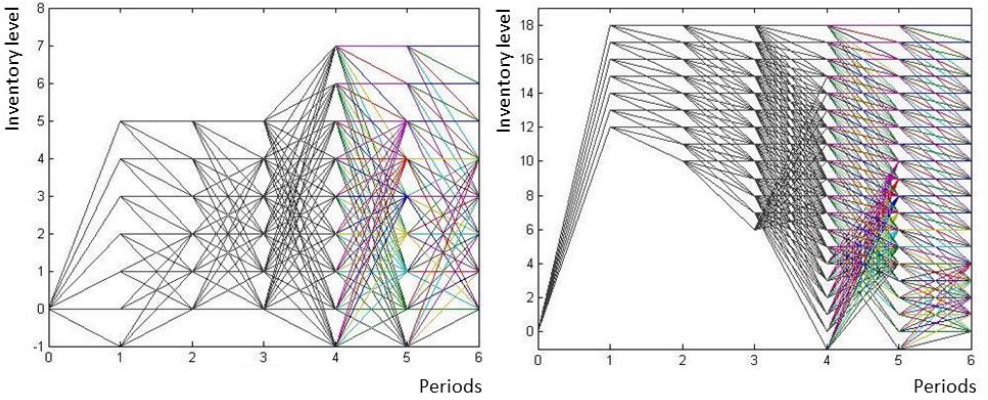


Fig. 3.3 Inventory development for all demand paths of the SDP solution with α -service level constraint for $k = 5$ (left) and $k = 50$ (right). Inventory levels below zero represent lost sales.

The SDP policies taught us that an optimal policy may consist of order-up-to levels for every period t , i.e. there is an order-up-to value S_t for each period such that the policy chooses as order quantity

$$Q_t(I) = (S_t - I)^+ \quad t = 1, \dots, T \quad (16)$$

Although such a policy does not necessarily lead to the lowest possible expected costs, we can look for the best values such that the α -service level constraint (10) is met. The optimal order-up-to level policy can then be compared to the SDP policy. The small instance allows us to perform a full enumeration to find the best feasible values for (S_1, \dots, S_T) . The enumeration procedure uses bounds for the order-up-to levels, being

$$S_1 \in \left[\hat{Q}_t, \sum_{t=1}^T 2\mu_t \right] \text{ and } S_t \in \left[0, \sum_{j=t}^T 2\mu_j \right] \quad t = 2, \dots, T \quad (17)$$

Systematically S_t is lowered and checked for feasibility of the service level constraint. The minimum expected total costs order-up-to levels are determined in Algorithm 3.1. The

number of possible combinations of the order-up-to levels within this range is large, about

$\prod_{t=1}^T \sum_{j=t}^T N_j$. For the case this are around 10^7 combinations. To compare, SDP evaluates in

each stage $t = 1, \dots, T$, for each value of I less than $\sum_{j=t}^T N_j$ possibilities for the value of Q .

Algorithm 3.1: Order-up-to $(\mu_1, \dots, \mu_T, \alpha)$: Determines S_t for $d_t \sim U\{0, 2\mu_t\}$ fulfilling (10) via (16)

- 1: Determine all possible demand paths for d_1, \dots, d_T
 - 2: Determine upper bounds on S_t
 - 3: Evaluate all combinations S_1, \dots, S_T within the bounds and keep the lowest cost solution
-

The generated order policy is listed in Table 3.3. The expected total costs $E(TC)$ of the SDP policy are respectively 12.7% and 19.0% higher than those of the optimal order-up-to policy. Fig. 3.4 shows the inventory development for all possible demand paths. Compared to the SDP solution, more out-of-stocks are allowed represented by a negative inventory level.

Table 3.3 Optimal order-up-to policy and its reached service levels (required $\alpha = 0.8$)

$k = 5; E(TC) = 32.79$							$k = 50; E(TC) = 108.37$						
Service	1.00	0.86	0.86	1.00	0.83	0.86	Service	1.00	1.00	1.00	0.996	0.90	0.80
\hat{Q}_t	5	2	4	7	5	4	\hat{Q}_t	5	2	4	7	5	4
I_{t-1}	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6	I_{t-1}	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6
0	6	0	3	8	4	3	0	18	0	0	7	0	0
1	0	0	2	7	3	2	1	0	0	0	6	0	0
2	0	0	1	6	2	1	2	0	0	0	5	0	0
3	0	0	0	5	1	0	3	0	0	0	4	0	0
4	0	0	0	4	0	0	4	0	0	0	3	0	0
5	0	0	0	3	0	0	5	0	0	0	2	0	0
6	0	0	0	2	0	0	6	0	0	0	1	0	0
7	0	0	0	1	0	0	7	0	0	0	0	0	0
:	0	0	0	0	0	0	:	0	0	0	0	0	0
10	0	0	0	0	0	0	18	0	0	0	0	0	0

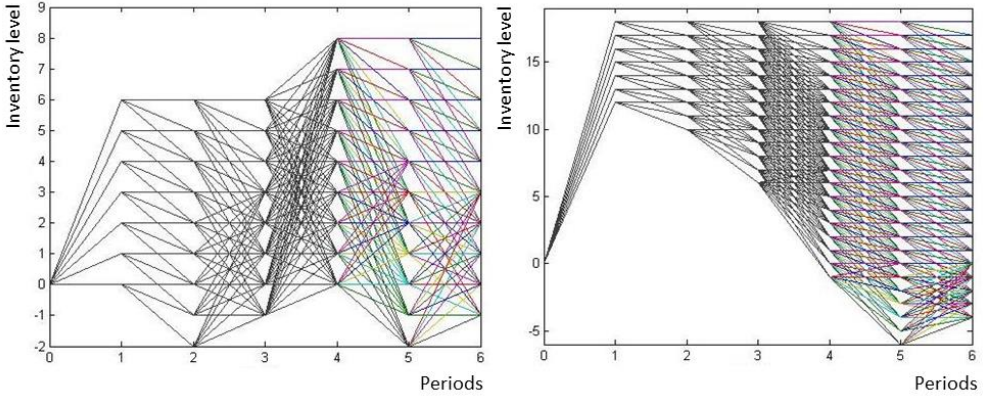


Fig. 3.4 Inventory development for all demand paths of the order-up-to policy with α -service level constraint for $k = 5$ (left) and $k = 50$ (right)

3.5 Fill rate constraint

Instead of an α -service level constraint we now consider a β -service level or fill rate constraint in the problem. The fill rate indicates that a predefined percentage of the demand per period has to be fulfilled from stock. Demand that cannot be fulfilled from stock is lost. Lost sales is defined as

$$X_t = (d_t - I_{t-1} - Q_t)^+ \quad t = 1, \dots, T \quad (18)$$

The fill rate constraint we consider is defined for each period

$$E(X_t) \leq (1 - \beta)\mu_t \quad t = 1, \dots, T \quad (19)$$

The fill rate constraint requires that the expected shortage is smaller or equal to a fraction $(1 - \beta)$ of the expected demand per period. For the SDP approach, now the feasible set $F(I)$ should take (19) into account.

Table 3.4 SDP order policy and its reached fill rates (required $\beta = 0.8$)

$k = 5; E(TC) = 32.30$							$k = 50; E(TC) = 122.92$						
Fill rate	0.86	1.00	0.94	0.83	0.87	0.92	Fill rate	1.00	1.00	1.00	0.99	0.98	0.97
I_{t-1}	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6	I_{t-1}	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6
0	4	5	3	5	4	3	0	17	15	15	14	9	3
1	0	4	2	4	3	2	1	0	14	14	13	8	2
2	0	0	1	3	2	1	2	0	0	13	12	7	1
3	0	0	0	2	1	0	3	0	0	0	11	6	0
4	0	0	0	1	0	0	4	0	0	0	10	0	0
5	0	0	0	0	0	0	5	0	0	0	0	0	0
:	0	0	0	0	0	0	:	0	0	0	0	0	0
7	0	0	0	0	0	0	17	0	0	0	0	0	0

This problem has been solved by an SDP approach implemented and evaluated in MATLAB. The results of the SDP model and the evaluated fill rates are listed in Table 3.4 and Fig. 3.5 shows all the demand paths in cases $k = 5$ and $k = 50$. Negative inventory levels represent out-of-stocks related to lost sales.

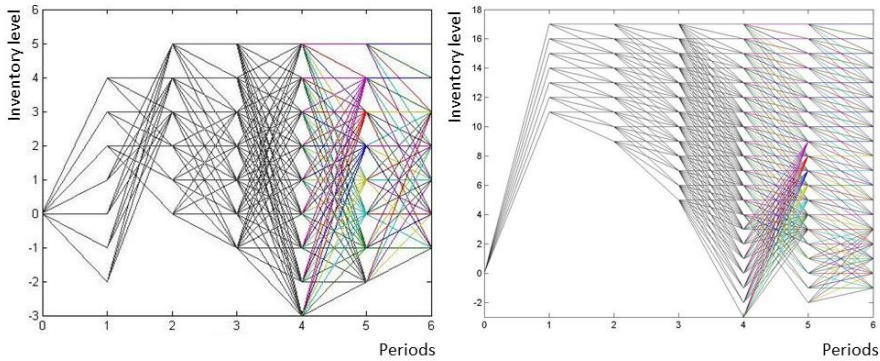


Fig. 3.5 Inventory development for all demand paths of the SDP solution with fill rate constraint for $k = 5$ (left) and $k = 50$ (right). Inventory levels below zero represent out-of-stocks (lost sales).

As can be observed, also with a fill rate constraint the order quantity $Q_t(I)$ follows an order-up-to level policy and when $k = 50$, for every period there is a reorder point, i.e. an inventory level above which no order takes place. The evaluated fill rates are in most periods well above the target fill rate of $\beta = 0.8$. As described in Section 3.4, from all possible starting inventory levels, SDP wants to meet the fill rate requirement. This results in an overall overachievement of the service level. The fill rate constraint we want to meet is described by (19). Instead, the SDP approach meets the conditional constraint

$$E(X_t | I_{t-1}) \leq (1 - \beta)\mu_t \quad t = 1, \dots, T. \quad (20)$$

Essentially the same algorithm as described in Section 3.4 can be used to find the optimal order-up-to values solution satisfying (19). The results are listed in Table 3.5 and shown in Fig. 3.6.

Table 3.5 Optimal order-up-to policy and reached fill rates (required $\beta = 0.8$)

$k=5$; $E(TC) = 30.03$							$k=50$; $E(TC) = 111.81$						
Fill rate	1.00	0.81	0.81	0.97	0.90	0.80	Fill rate	1.00	1.00	1.00	0.999	0.95	0.80
I_{t-1}	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6	I_{t-1}	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6
0	6	0	2	7	4	2	0	18	0	0	7	3	0
1	0	0	1	6	3	1	1	0	0	0	6	2	0
2	0	0	0	5	2	0	2	0	0	0	5	1	0
3	0	0	0	4	1	0	3	0	0	0	4	0	0
4	0	0	0	3	0	0	4	0	0	0	3	0	0
5	0	0	0	2	0	0	5	0	0	0	2	0	0
6	0	0	0	1	0	0	6	0	0	0	1	0	0
7	0	0	0	0	0	0	7	0	0	0	0	0	0
8	0	0	0	0	0	0	:	0	0	0	0	0	0
9	0	0	0	0	0	0	18	0	0	0	0	0	0

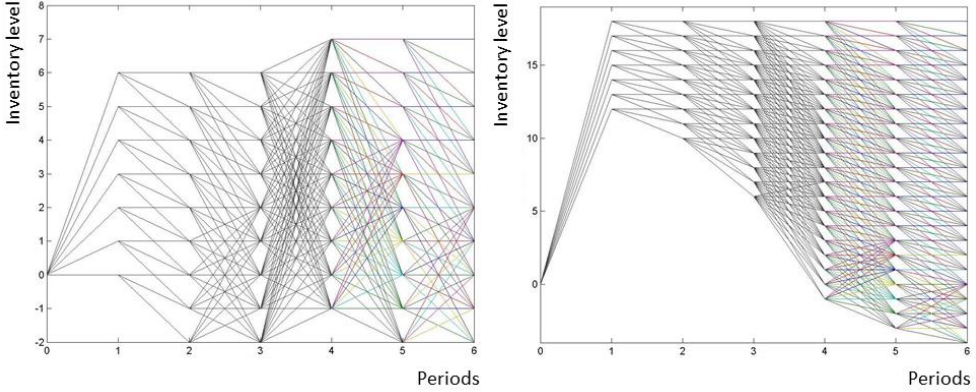


Fig. 3.6 Inventory development of all demand paths of the order-up-to policy with fill rate constraint for $k = 5$ (left) and $k = 50$ (right)

The expected total costs $E(TC)$ of the SDP policy are respectively 7.6% and 9.9% higher than of the optimal order-up-to policy. In this example, with a fill rate constraint the SDP solution is closer to the optimal order-up-to policy than with an α -service level constraint.

3.6 Continuous demand

The use of the discrete Uniform distribution to model the demand enabled us to calculate order policies with SDP and to evaluate the policies with all possible demand paths. However, the discrete Uniform distribution is not very realistic to model demand. In cases of small demand numbers, the Gamma distribution can be suitable (Burgin 1975). We consider the Gamma distribution $d \sim \Gamma(K, \vartheta)$, where K is the shape parameter and ϑ is the scale parameter. The expected value of the Gamma distribution is $E[d] = \vartheta K$. We assume $\vartheta = 1$, then $K = \mu$. As before, the expected demand $\mu_t = 3, 1, 2, 4, 3, 2$. Let now $d_t \sim \Gamma_t(\mu_b, 1)$ with G_t as corresponding cumulative distribution function (cdf).

To apply SDP, demand has to be discretized. A common method is to discretize the space of possible outcomes of the stochastic demand δ_{ij} by using quantiles of the Gamma distribution. An equidistant grid is taken over the range $[0, 1]$ in N steps, each with probability $p = 1/N$. The discrete outcomes are $\delta_{ij} = G_t^{-1}(j \times p)$, $j = 1, \dots, N$ (Van Dijk et al. 2014). Each outcome has the same probability of occurrence and the outcome space is truncated. Because of the skewness of the Gamma distribution, there are more outcomes with demand smaller than or around μ_b than greater than μ_b . In Fig. 3.7 this principle is illustrated for $N = 10$ steps with $p = 0.1$.




Fig. 3.7 Discretisation of Gamma distribution with $\mu_t = 4$ and $\rho = 0.1$

After discretisation, value function (7) can be reformulated into

$$V_t(I) = \min_{Q \in F_t(I)} \left[g(Q) + h \cdot \rho \sum_{j=1}^{N_t} (I + Q - \delta_{tj})^+ + \rho \sum_{j=1}^{N_t} V_{t+1}(I + Q - \delta_{tj})^+ \right] \quad (21)$$

with feasible area

$$F_t(I) = [0, (\delta_{tN_t} + \dots + \delta_{tN_T} - I)^+]. \quad (22)$$

The feasible set $F_t(I)$ should take the fill rate constraint (19) into account, so at least in period 1, the actual feasible set is smaller, but there is no easy analytical expression for the lower bound in case of a Gamma-distributed fill rate. The number of possible demand paths depends on the number of steps in the grid. In this case we use a grid of $N = 100$ steps in every period, so there are $100^T = 10^6$ possible paths. SDP still generates a solution within a second, because each stage, only 100 outcomes have to be considered, but the solution cannot be evaluated anymore by generating all paths, it has to be simulated. Notice that the demand outcomes lead to inventory levels that do not have an integer (grid) value. This means that the value of $V_{t+1}(I + Q - \delta_{tj})^+$ in (21) has to be found by interpolation for each demand outcome δ_{tj} in order to get the expected value for an order choice Q .

Table 3.6 SDP order quantities given inventory level, Gamma-distributed demand, simulated costs and reached fill rates (required $\beta = 0.8$)

$k = 5; V_1(0) = 31.82; \text{SimAvg(TC)} = 33.88$							$k = 50; V_1(0) = 120.59; \text{SimAvg(TC)} = 123.99$						
Fill rate	0.95	0.95	0.86	0.81	0.94	0.85	Fill rate	1.00	1.00	1.00	0.99	0.98	0.94
I_{t-1}	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6	I_{t-1}	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6
0	5.12	3.72	2.50	4.06	5.01	2.50	0	16.78	14.66	14.34	13.01	8.80	2.50
1	0	2.71	1.51	3.06	4.00	1.50	1	0	13.66	13.34	12.00	7.79	1.50
2	0	0	0.50	2.05	3.01	0.50	2	0	0	12.34	11.00	6.80	0.50
3	0	0	0	1.04	2.01	0	3	0	0	0	9.99	5.80	0
4	0	0	0	0.04	0	0	4	0	0	0	8.99	0	0
5	0	0	0	0	0	0	5	0	0	0	0	0	0
6	0	0	0	0	0	0	:	0	0	0	0	0	0
7	0	0	0	0	0	0	17	0	0	0	0	0	0

For the final evaluation of the policy, 20000 random Gamma distributed demand paths are simulated. In Table 3.6 the SDP order policies are listed, together with the results of the simulation, for $k = 5$ and $k = 50$. Fig. 3.8 shows the demand paths for both cases. Due to the continuous distribution of the demand, the optimal order quantities do no longer have an integer value and are found by a one-dimensional nonlinear optimization over the interval $F_t(I)$ for a set of integer valued grid points of the inventory. The order quantities $Q_t(I)$ follow an order-up-to level policy, with a reorder point in some periods. In case of $k = 5$, only in period 4 the fill rate is close to the requirement. In the other periods, and in case of $k = 50$, an overachievement of the fill rate can be observed, as in the discrete model in Section 3.5.

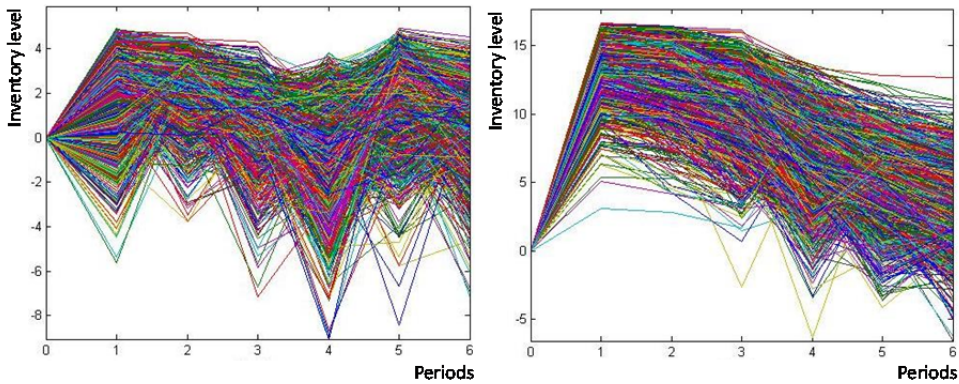


Fig. 3.8 Inventory development for sampled demand from Gamma distribution of the SDP policy with fill rate constraint, $k = 5$ (left) and $k = 50$ (right).

Algorithm 3.1 is based on enumerating integer values for order-up-to levels S_t , whereas in the continuous case they can take a continuous value. Moreover, the number of demand paths is not finite, so we rely on sample paths. The optimization of S_t was based on nonlinear optimization using the `fmincon` solver of `MATLAB` and using 50000 sample paths to evaluate the costs and to estimate the fill rate. Table 3.7 gives the resulting solution of the order-up-to levels.

Table 3.7 Optimal order-up-to level quantities for inventory level, Gamma-distributed demand, simulated costs and reached fill rates (required $\beta = 0.8$)

$k = 5$; SimAvg(TC) = 30.35							$k = 50$; SimAvg(TC) = 105.54						
Fill rate	0.96	0.80	0.80	0.80	0.98	0.81	Fill rate	1.00	1.00	1.00	1.00	0.93	0.80
I_{t-1}	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6	I_{t-1}	Q_1	Q_2	Q_3	Q_4	Q_5	Q_6
0	5.53	0	2.03	3.92	6.13	0.23	0	17.67	0.66	3.06	0	0	0
1	0	0	1.03	2.92	5.13	0	1	0	0	2.06	0	0	0
2	0	0	0.03	1.92	4.13	0	2	0	0	1.06	0	0	0
3	0	0	0	0.92	3.13	0	3	0	0	0.06	0	0	0
4	0	0	0	0	2.13	0	4	0	0	0	0	0	0
5	0	0	0	0	1.13	0	5	0	0	0	0	0	0
6	0	0	0	0	0.13	0	:	0	0	0	0	0	0
7	0	0	0	0	0	0	18	0	0	0	0	0	0

The simulated total costs $\text{SimAvg}(\text{TC})$ of the SDP policy are respectively 11.6% and 17.5% higher than of the order-up-to level policy. The reached fill rates of the order-up-to level policy are close to the requirements. This specific example proves that the overachievement of the fill rate constraint by SDP due to the conditional treatment, also occurs for continuously distributed demand.

3.7 Conclusions

By means of a stylized example of an inventory management problem that is relevant for practice we study whether SDP is a suitable method if a minimal service level constraint applies and demand is non-stationary. The example has a fixed time horizon is six periods. We started to solve a deterministic problem, where the variable demand equals the expected demand in the stochastic cases. Introducing uncertainty through a Uniform distribution and requiring that all demand has to be fulfilled gives a 75% cost increase. Subsequently we introduced an α -service level constraint and a fill rate constraint per period, being minimal service level constraints.

Evaluation of the SDP order policies shows that SDP generates order policies that overachieve the service level requirements. This is due to the conditional behaviour of the service level constraints. From all possible starting inventory levels, no matter how small the chance of occurrence, SDP has to meet the service level requirement per period. We could compare the SDP policy with an order-up-to level policy, which in principle has less choice for the order quantity. The corresponding order-up-to levels were generated by full enumeration meeting the service level requirement, resulting in a policy with lower costs and service levels that are at the end of the replenishment cycle close to the required service level.

Finally, continuously distributed demand is modelled with a more realistic Gamma distribution and a fill rate constraint. Compared to the Uniform distribution, this results in lower order quantities and lower costs, but also in an overachievement of the service level requirement. Future research may aim at investigating the effect of the overachievement of SDP of service levels for other demand distributions. Due to the conditional fulfilment of the service levels, we observe this overachievement. However this effect may be less for other type of distributions and therefore SDP might generate acceptable solutions.

To conclude: SDP generates order policies that meet a conditional service level constraint. Given a minimal expected service level requirement, the resulting policy is not necessarily optimal.

Chapter 4

Comparing order policies and solution methods

Based on:

Pauls-Worm, K.G.J., Hendrix, E.M.T., Rossi, R., 2016. *Inventory control for a non-stationary demand perishable product: comparing policies*. Submitted to an international journal.

Abstract

This chapter studies inventory control for a perishable product with a maximum fixed shelf life in a periodic review system. A Stochastic Programming (SP) problem is formulated to model the practical production planning problem over a finite horizon. Perishability, non-stationary demand, fixed ordering cost and a service level requirement make this problem complex. Furthermore, various inventory policies can be defined to handle the inventory control problem, from easy to implement to more complex decision rules. We study three different policies: a) the production timing is fixed in advance combined with an order-up-to level, b) the production timing is fixed in advance and the production quantity takes the age-distribution into account and c) every period is decided to produce or not, and how much, depending on the age-distribution of the items in stock. Chapter 2 used an MILP approximation to derive values for policy a). In this chapter we use a computational method based on the so-called Smoothed Monte Carlo method with sampled demand to optimize values for such a system. The resulting MINLP approach uses enumeration, bounding and iterative nonlinear optimisation. Moreover, we present a sample based approach to determine the order quantities for policy b). Flexible policy c) is derived by SDP. All policies are compared. In most tested scenarios the expected total costs of the policies are very close and the MILP values for policy a) provide a cost efficient and easy to implement solution. In situations where the cost structure is such that the replenishment cycle lengths are highly varying, waste will occur during the replenishment cycle and the parameters generated by MINLP may be more suitable. Only in situations of a high service level and a coefficient of variation of 0.25 or more, the age-distribution is important. Furthermore, the flexible policy is only appropriate when setup cost is low, a coefficient of variation of 0.33 or more and a high service level of 98%.

4.1 Introduction

The demand of perishable food products is always uncertain and often also fluctuating. Partly those demand fluctuations are due to promotions, so they are planned. In production and inventory control of perishable products, one has to deal with this non-stationary demand by varying production and inventory per period.

In this chapter we consider perishable products that are processed and get a best-before-date. Examples are packed cheese, cut and packed lettuce, yoghurt, etc. Those products have a fixed maximum shelf life. Producers and retail organisations have arrangements about the necessary available shelf life for the consumer, this means each actor in the supply chain has a limited number of periods for its logistic activities and planning to preserve enough shelf life for the other actor(s) and the consumer. We can regard the available periods for logistic activities as the maximum (internal) shelf life. After the maximum shelf life the product cannot be used anymore for the intended purpose and is considered waste. An α -service level requirement has to be met, i.e. the probability to be out of stock should be smaller than $1 - \alpha$. For every period the producer has to decide whether or not to order and how much, considering a fixed ordering cost, holding cost and disposal cost. This results in replenishment cycle lengths of a varying number of periods. The producer has control over the issuing of products, so in order to minimise waste, a First-in-First-Out (FIFO) policy is used. Excess demand is backlogged.

We modelled this complex practical problem for a fixed horizon of T periods as a Stochastic Programming problem published in (Pauls-Worm et al., 2014) (Chapter 2). Based on this model various inventory policies can be defined to handle the inventory control problem. We study three different policies a) the production timing is fixed in advance combined with an order-up-to level, b) the production timing is fixed in advance and the production quantity takes the age-distribution is taken into account and c) every period is decided to produce or not, and how much, depending on the age-distribution of the items in stock. In inventory literature, in general the replenishment cycle R is the length of the period, or the number of periods, for which the order-up-to level S and the order quantity Q is meant. For stationary demand, the replenishment cycle is fixed, but for non-stationary demand the optimal replenishment cycle R_t may depend on order moment t . In our case, replenishment cycle R_t for order moment t , where $Y_t = 1$, is the number of periods (days, weeks,...) for which the order aims to cover demand. If no order takes place $Y_t = 0$. In our policies we want to determine whether we order or not at the beginning of a period, therefore we use the notation Y instead of R . Let X denote the inventory age-distribution at the beginning of a period. Given the introduced symbols, the described policies a), b) and c) will be called YS , $YQ(X)$ and $Q(X)$ policy, respectively.

In the practical planning of a food producer it is convenient to know beforehand in which periods to produce, following an order timing Y and order-up-to level S . This YS policy is easy to implement in practice. Chapter 2 (Pauls-Worm et al., 2014) used an MILP approximation to derive values for a YS system. In situations where the cost structure is such that the replenishment cycle lengths are highly varying, waste will occur during the replenishment cycle. Consequently, the developed method does not always fulfil the

required service level. In this study we use a computational method based on the so-called Smoothed Monte Carlo method with sampled demand to optimise values for a YS system. The resulting MINLP approach uses enumeration, bounding and iterative nonlinear optimisation.

In a YS policy, only the expected waste during a replenishment cycle is considered in the order-up-to level. The order quantity is determined by the order-up-to level minus the inventory in stock. It would be more cost-efficient to consider the age-distribution of all items in stock in determining the order quantity. We present a sample based $YQ(X)$ policy in which the age-distribution is taken into account. Finally, we can consider a more flexible $Q(X)$ policy derived by SDP. In the $Q(X)$ policy, at the start of every period the order quantity is determined, based on the age-distribution of the items in stock.

The research question of this chapter is in which cases which policy is most suitable. In Section 4.2 we present a literature review on inventory control policies in case of non-stationary demand for a perishable item with a fixed shelf life. In Section 4.3 the SP model of the practical problem is given. Section 4.4 explains some characteristics of a solution for a YS or $YQ(X)$ policy. Section 4.5 describes an MINLP approach to determine parameters for a YS policy. The $YQ(X)$ policy is depicted in Section 4.6 and Section 4.7 discusses the flexible $Q(X)$ policy. Section 4.8 compares the policies on effectiveness and efficiency and Section 4.9 concludes.

4.2 Literature

To answer the question which order policy is most suitable in which situation, we first focus in the literature on periodic review order policies for non-stationary demand. In the second part we take perishability with a fixed shelf life into consideration.

Bookbinder and Tan (1988) distinguish three strategies to deal with ordering of non-perishable products with non-stationary demand in a periodic review. Firstly, the static uncertainty strategy, where the timing and size of the orders must be known at the beginning of the time horizon, results in what we call a YQ policy. This policy is appropriate when there is considerable lead time (Pauls-Worm et al., 2015) (Chapter 6). Secondly, the static-dynamic uncertainty strategy, where timing of the orders is known at the beginning of the time horizon, but the order quantity may be adapted in response of the inventory levels observed during the time horizon. We call this a YS policy. In a heuristic approach Bookbinder and Tan (1988) split the problem in two stages. The first stage determines the timing of the order, the second the quantity. Tarim and Kingsman (2004) considered the Bookbinder and Tan approach as a basis for the formulation of a MILP model for non-stationary stochastic demand for the simultaneous determination of the timing and size of the replenishment orders. Tarim and Kingsman's approach provides an optimal solution. Thirdly, the dynamic uncertainty strategy where the order quantity is decided at the beginning of every period, results in what we call a $Q(X)$ policy. This type of order policy is discussed already in the 1960s. Karlin (1960a) shows that a critical number policy is optimal, were the critical numbers are a reorder level s_t and an order-up-to level S_t resulting in an

(Rs_tS_t) policy. Karlin (1960b) and Veinott Jr (1963; 1965) developed optimal myopic policies for certain cases. Morton (1978) shows that near myopic bounds are close to optimal, under the assumption of disposal of excess stock. Morton and Pentico (1995) derive near-myopic bounds for the more general case. Zipkin (1989) developed optimal critical number policies for a cyclic demand pattern. He shows that the critical numbers in the optimal policy smooth the fluctuation in the demand data. This “wait-and-see” approach in the critical number policies following the dynamic uncertainty strategy could require an order with setup cost in almost every period. This might be undesirable for the production planning of a company, but in case of large setup cost relative to the holding cost, this is also not optimal (Bookbinder and Tan, 1988).

Order policies for perishable products with a fixed lifetime are reviewed by Nahmias (1982), Goyal and Giri (2001), and Karaesmen et al. (2011). Almost all papers assume stationary demand. Fries (1975) shows that with a maximum shelf life of $M \geq 2$, neither an RS nor an RsS policy is optimal. Nahmias (1975) and Fries (1975) observe that in general an optimal order policy for perishable products with a fixed life time should take the age-distribution of the products in stock into account. Even when all perishable items are of the same age, base stock (order-up-to) policies are not optimal, as argued by Tekin et al. (2001) and Haijema et al. (2007). Some papers, e.g. (Haijema et al., 2007), (Broekmeulen and Van Donselaar, 2009), (Minner and Transchel, 2010) assume a cyclic demand pattern, with a weekly demand pattern per day, but stationary expected demand per week. They assume negligible setup cost and follow a dynamic uncertainty strategy, which might not be optimal in case of fixed setup cost. To our knowledge, non-stationary demand which is not cyclic is hardly studied for perishable products, and also the static-dynamic uncertainty strategy to establish an order policy is little exposed. The static-dynamic uncertainty strategy is appropriate for planning purposes.

Pauls-Worm et al. (2014) formulated an SP inventory model in Chapter 2 that minimises the expected total costs, including setup cost, unit procurement cost, holding cost and cost of waste, for a perishable product with non-stationary stochastic demand with an α -service level constraint under a FIFO issuing policy. In (Pauls-Worm et al., 2014) (Chapter 2) and (Hendrix et al., 2015) the SP model is solved for the static-dynamic uncertainty strategy.

4.3 Stochastic Programming model

The SP model formulated in Chapter 2 (Pauls-Worm et al., 2014) is described below. Due to the stochastic demand d_t , the inventory levels I_{bt} (except the starting inventory I_{b0}) and the order quantities Q_t are random variables. We assume a lead time of zero. The probability in the chance constraints is notated as $P(\cdot)$, and $E(\cdot)$ denotes the expected value operator. We use $(x)^+$ to express $\max \{x, 0\}$. In Table 4.1 the list of symbols can be found.

Table 4.1 List of symbols

Indices	
t	period index, $t = 1, \dots, T$ with T the time horizon
b	age index, $b = 1, \dots, M$ with M the fixed maximum (internal) shelf life
Data	
d_t	Normally distributed demand with expectation $\mu_t > 0$ and variance $(CV \times \mu_t)^2$ where CV is a given coefficient of variation
k	fixed ordering or setup cost, $k > 0$
c	unit procurement cost, $c > 0$
h	unit holding cost, for items that are carried over from one period to the next, $h > 0$
w	unit disposal cost ($w > 0$) or salvage value ($w < 0$) for items becoming waste
α	required service level, $0 < \alpha < 1$
Variables	
Q_t	ordered and delivered quantity at the beginning of period t
I_{bt}	inventory level of items with age b at the end of period t , initial inventory fixed $I_{b0} = 0$, $I_{1t} \in \mathbb{R}$, $I_{bt} \geq 0$ for $b = 2, \dots, M$. Inventory of age M at the end of period t is considered waste

The expected total costs over the time horizon T are minimized.

$$\text{Min } E(TC) = E \left(\sum_{t=1}^T \left(g(Q_t) + h \sum_{b=1}^{M-1} I_{bt}^+ + w I_{Mt} \right) \right) = \sum_{t=1}^T E \left(g(Q_t) + h \sum_{b=1}^{M-1} I_{bt}^+ + w I_{Mt} \right) \quad (1)$$

where the procurement costs are given by the function

$$g(x) = k + cx \text{ if } x > 0, \text{ and } g(0) = 0. \quad (2)$$

The chance constraint requiring the α -service level is expressed by

$$P(I_{1t} \geq 0) \geq \alpha \quad t = 1, \dots, T \quad (3)$$

meaning that the probability of not being out-of-stock at the end of period t is greater than or equal to α . The probability of a stock-out is $1 - \alpha$. Because of a FIFO issuing policy, the inventory levels of the older items are zero in case of a shortage, so only the inventory of the freshest items can be negative. The inventory dynamics for the FIFO issuing is described by

$$I_{bt} = \left(I_{b-1, t-1} - \left(d_t - \sum_{j=b}^{M-1} I_{j, t-1} \right)^+ \right)^+ \quad t = 1, \dots, T; \quad b = 2, \dots, M \quad (4)$$

and for the freshest items

$$I_{1t} = Q_t - \left(d_t - \sum_{b=1}^{M-1} I_{b, t-1} \right)^+ \quad t = 1, \dots, T \quad (5)$$

Items of age M become waste at the end of the period and cannot be used in the next period.

The inventory of the freshest items can be negative. The unmet demand will be backlogged. The inventory at the beginning of a period is defined by

$$X = (I_{1,t-1}, \dots, I_{M-1,t-1}) \quad (6)$$

As described in Section 4.1, several order policies can be defined as solutions of the SP model. One way is to define order-up-to levels S_b . The decision maker replenishes in period t the inventory up-to the level S_b where the order quantity Q_t is defined by

$$Q_t(X) = \left(S_t - \sum_{b=1}^{M-1} X_{bt} \right)^+ \quad (7)$$

In Section 4.4 we discuss characteristics of policy solutions according to the static-dynamic uncertainty strategy for this SP model. In Section 4.5 we develop a YS policy and in Section 4.6 a $YQ(X)$ policy, both following a static-dynamic uncertainty strategy. In Section 4.7 we discuss a more flexible $Q(X)$ policy according to the dynamic uncertainty strategy to obtain parameters for the SP model.

4.4 Characteristics of a solution to the static-dynamic uncertainty strategy

First we focus on the characteristics of feasible solutions. In Section 4.4.1, the concept of replenishment cycles and the timing of replenishments are discussed. In Section 4.4.2, we show in which cases the so-called basic order-up-to level is the optimal quantity. For the other order moments we study the mathematical implications of estimating the service level by a Monte Carlo sampling approach in Section 4.4.3.

4.4.1 Replenishment cycles and limits on timing vector Y

A replenishment cycle is the number of periods R the order quantity Q_t aims to fulfil. For non-stationary demand, replenishment cycle length R can depend on order moment t . In terms of the order timing vector Y , it means that $Y_t = 1$ and $Y_{t+R+1} = 1$ and no orders take place in between. In case of perishable items with maximum shelf life M , the replenishment cycle cannot be longer than the shelf life M , so $R \leq M$.

Lemma 1. *Let Y be an order timing vector of the SP model, i.e. $Y_t = 0 \Rightarrow Q_t = 0$. Y provides an infeasible solution of the SP model, if it contains more than $M - 1$ consecutive zeros.*

F_T is defined as the set of all feasible order timing vectors Y of length T . The number of elements $|F_T|$ of the set F_T of feasible order timings of T periods and a shelf life of $M < T - 1$ follows the recursive rule $|F_{T+1}| = 2|F_T| - |F_{T-M}|$ with the initial terms $|F_t| = 2^{t-1}$ for $t < M + 1$ and $|F_{M+1}| = 2^M - 1$ as shown in (Alcoba et al., 2015). F_T is exponential in the horizon T . However, in a practical situation, as explained in Chapter 2 (Pauls-Worm et al., 2014), it is sufficient to plan ahead for $T = 12$ periods.

4.4.2 Basic order-up-to level and optimal order quantities

For a certain replenishment cycle length $R = 1, \dots, M$ we can define a basic order-up-to level \hat{S}_{Rt} as the inventory that should be available at the beginning of period t to cover demand of R periods.

Definition 1. Let $d_t + \dots + d_{t+R-1}$ be the stochastic demand during the replenishment cycle of length R with cumulative distribution function (cdf) G_{Rt} . The basic order-up-to level \hat{S}_{Rt} with probability α to fulfil demand is defined by $G(\hat{S}_{Rt}) = \alpha$ such that $\hat{S}_{Rt} = G_{Rt}^{-1}(\alpha)$.

For some replenishment cycles, to have the basic order-up-to level in stock at the beginning of the period may be not enough, so $\sum_{b=1}^{M-1} X_{bt} + Q_t \geq \hat{S}_{Rt}$, because products in stock can become waste during the replenishment cycle. Nevertheless, for some replenishment cycles, the basic order-up-to level is sufficient and specifies the optimal order quantity. First consider an order moment following a replenishment cycle of the length of maximum shelf life M .

Lemma 2. Let Y be an order timing vector of the SP model with corresponding cycle length R and X defined by (6). For an order moment t having $Y_{t-M} = 1$, $R = M$, the optimal order quantity is $Q_t = \hat{S}_{Rt}$.

Proof. Minimisation of the costs leads to a value of Q_t as low as possible. Constraints (4) and (5) define that after M periods no (non-perished) inventory is left over from order Q_{t-M} , so $I_{b,t-1} = 0$ for $b = 1, \dots, M-1$. To fulfil chance constraint (3), the order quantity should fulfil $Q_t \geq \hat{S}_{Rt}$. Minimising its value implies $Q_t = \hat{S}_{Rt}$. \square

Second, there may be replenishment cycles of just one period where during the cycle no products can become waste.

Lemma 3. Let Y be an order timing vector of the SP model and X defined by (6). For an order moment t having $Y_{t+1} = Y_t = 1$ the optimal order quantity is $Q_t = \hat{S}_{1t} - \sum_{b=1}^{M-1} X_{bt}$.

Proof. For one period demand, chance constraint (3) translates to $P\left(Q_t + \sum_{b=1}^{M-1} X_{bt} \geq d_t\right) = \alpha$

$\Rightarrow Q_t \geq \hat{S}_{1t} - \sum_{b=1}^{M-1} X_{bt}$. Minimising its value implies $Q_t = \hat{S}_{1t} - \sum_{b=1}^{M-1} X_{bt}$. \square

Independently of the order timing, the best order quantity at a negative stock level always has an order-up-to character.

Lemma 4. Let Y be an order timing vector of the SP model and X defined by (6), $Y_t = 1$ with replenishment cycle length R . If $X_{1t} \leq 0$, the optimal order quantity is $Q_t = \hat{S}_{Rt} - X_{1t}$.

Proof. No old stock is available that can perish during the replenishment cycle. The chance constraint (3) translates to $Q_t + X_{1t} \geq \hat{S}_{Rt}$. Minimising the value of Q_t implies $Q_t = \hat{S}_{Rt} - X_{1t}$. \square

4.4.3 Monte Carlo estimation of the service level

With the theoretical properties in Section 4.4.1 and 4.4.2 we can determine the optimal order quantities for some of the order moments of a certain order timing vector Y . For the other order moments, the question is how to find the order quantities, considering the chance constraint. A way to deal with that is using samples of the demand series. Let for an order period t , \mathbf{d} be the stochastic demand vector $(\mathbf{d}_t, \dots, \mathbf{d}_{t+R-1})$ during replenishment cycle length R , X the starting inventory and Q the order quantity. Let $f(Q, X, \mathbf{d}) = I_{1,t+R-1}$ define the end inventory of items with an age of one period given a realisation \mathbf{d} of \mathbf{d} following the inventory dynamics with possible perishing according to (4) and (5). Consider the indicator function $\delta : \mathbb{R} \times \mathbb{R}^R \rightarrow \{0, 1\}$

$$\delta(Q, \mathbf{d}) = \begin{cases} 1 & \text{if } f(Q, X, \mathbf{d}) \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

that translates the service level in constraint (3) for period $t + R - 1$ to

$$a(Q) = P(I_{1,t+R-1} \geq 0) = E_{\mathbf{d}} \delta(Q, \mathbf{d}). \quad (9)$$

The generic concept of its translation is given in handbooks like (Jacod and Protter, 2004). Using sample paths $\mathbf{d}_1, \dots, \mathbf{d}_N$ of \mathbf{d} to estimate a probability like (9) was called by Von Neumann the Monte Carlo method. The idea is that given N sample paths $\mathbf{d}_1, \dots, \mathbf{d}_N$ of \mathbf{d} , the probability (service level) (9) is estimated by

$$\hat{a}(Q) = \frac{1}{N} \sum_{r=1}^N \delta(Q, \mathbf{d}_r). \quad (10)$$

As is known from handbooks on statistics (e.g. (Lyman Ott and Longnecker, 2001)), considering a set of independent random samples \mathbf{d}_r provides the unbiased estimator (10) of $a(Q)$ with standard deviation

$$\sigma(\hat{a}(Q)) = \sqrt{\frac{1}{N} (a(Q) - a(Q)^2)}. \quad (11)$$

The latter is of importance in Monte Carlo approaches to set the number of samples for a desired probabilistic accuracy. Following the idea that the binomial distribution is practically normal for a large number of samples, a rule of thumb is to have an accuracy of 2σ . For the particular application aiming at $\alpha = 90\%, 95\%, 98\%$, a sample size of $N = 5000$ gives a rule of thumb accuracy of about 0.005 of the estimator $\hat{a}(Q)$.

The next question is how to use the theoretical findings and the estimation method of Section 4.4 to find policies where the order timing Y is provided to the decision maker. We want to find the best order timing and a way to deal with the order quantity.

4.5 YS policy: sample based SMC-MINLP approach

We consider a static-dynamic uncertainty YS policy where the decision maker is provided an order timing vector Y , i.e. $Y_t = 0 \Rightarrow Q_t = 0$. The aim is to generate values for S_t such that the α -service level constraints are fulfilled for all instances and expected costs are minimised.

Lemmas 2 and 3 are helpful to define the order-up-to level $S_t = \hat{S}_{Rt}$ for specific moments. Sample-based estimation can be used for the service level in each period as presented in Section 4.4.3. Specifically, for the YS policy, one can write the problem of finding the (discrete) timing Y and (continuous) order-up-to levels S as a Monte Carlo based Mixed Integer Linear Programming (MC-MILP) model. Such a model is given in Section 4.5.1. It is known from literature that the corresponding sample-based model called MC-MILP for most instances cannot be solved in reasonable time, e.g. (Rijkema et al., 2013). Therefore, in Section 4.5.2 we investigate the possibility to use an equivalent MINLP model based on the Smoothed Monte Carlo method, see (Hendrix and Olieman, 2008). A specific algorithm is designed that uses enumeration and bounding for the integer part Y of the problem leaving us with iteratively solving an NLP problem in the continuous variables S . This policy is earlier described in (Hendrix et al., 2015).

4.5.1 MC-MILP optimisation of the YS policy

The sample-based approach for the YS policy can be handled by adding to the SP model a sample index $r = 1, \dots, N$ to the variables, I_{btr} and Q_{tr} such that one has replicas of the same variables that describe the actions of the model under each sample r . Furthermore, for the chance constraints one adds a binary variable $\delta_{tr} \in \{0, 1\}$ representing the indicator value that specifies whether demand is fulfilled in period t in sample r

$$-I_{1tr} \leq m_t(1 - \delta_{tr}) \quad r = 1, \dots, N; \quad t = 1, \dots, T \quad (12)$$

where m_t is an upper bound on the value of the out of stock $-I_{1t}$. This defines a function $\hat{\alpha}_t(S): \mathbb{R}^T \rightarrow \{0, \frac{1}{N}, \frac{2}{N}, \dots, 1\}$ representing the reached service level under the set of samples.

The corresponding chance constraints are

$$\hat{\alpha}_t(S) := \frac{1}{N} \sum_{r=1}^N \delta_{tr} \geq \alpha \quad t = 1, \dots, T \quad (13)$$

The objective (1) is extended towards

$$\text{Min } \frac{1}{N} \sum_{t=1}^T \left(kY_t + \sum_{r=1}^N \left(cQ_{tr} + h \sum_{b=1}^{M-1} I_{btr}^+ + wI_{Mtr} \right) \right) \quad (14)$$

with order quantity

$$Q_{tr} = \left(S_t - \sum_{b=1}^{M-1} I_{b,t-1,r} \right)^+ \quad r = 1, \dots, N; \quad t = 1, \dots, T \quad (15)$$

and the conventional order relation

$$S_t \leq \mathcal{M} Y_t \quad t = 1, \dots, T \quad (16)$$

with a big- \mathcal{M} value. The constraints (4) and (5) are extended to each sample

$$I_{btr} = \left(I_{b-1,t-1,r} - \left(d_{tr} - \sum_{j=b}^{M-1} I_{j,t-1,r} \right)^+ \right)^+ \quad r = 1, \dots, N; \quad t = 1, \dots, T; \quad b = 2, \dots, M \quad (17)$$

and for the freshest items

$$I_{1tr} = Q_{tr} - \left(d_{tr} - \sum_{b=1}^{M-1} I_{b,t-1,r} \right)^+ \quad r = 1, \dots, N; \quad t = 1, \dots, T \quad (18)$$

Notice that the values that we intend to find, i.e. Y_t and S_t are independent of the sample r and the other variables that describe the simulation or evaluation part Q_{tr} , I_{btr} and δ_{tr} depend on the sample.

Solving the MC-MILP model is in most cases practically impossible due to the large number of binary variables δ and many solutions δ that represent the same obtained service levels $a(S)$. The number of samples $N = 5000$ mentioned in Section 5.1, implies defining for each period $N = 5000$ binary variables δ_{tr} . Instead, we will investigate a smoothed Monte Carlo approach as suggested in (Hendrix and Olieman, 2008) to estimate the service levels in the MC-MILP model.

4.5.2 Smoothed Monte Carlo MINLP approach to the YS policy

Consider the MC-MILP problem from the point of view of an NLP problem in the continuous variables S when order timing Y is given. The function $\hat{a}_t(S): \mathbb{R}^T \rightarrow \{0, \frac{1}{N}, \frac{2}{N}, \dots, 1\}$ in (13) and objective (14) are evaluated by using N sample paths following the dynamics (16), (17) and (18). The difficulty of applying an NLP solver for this problem is that (13) is piecewise constant, i.e. changing the values of S a bit may not change the evaluated value of $\hat{a}_t(S)$. (Hendrix and Olieman, 2008) show that one can make the reached service level practically a continuous function by following the MC smoothing approach. Let $z_{rt} = \sum_{b=1}^M I_{btr}$

represent the total amount of product left over at the end of period t in sample r . One can measure how close $\hat{a}_t(S)$ is to change value by the value of the least nonnegative total inventory $p_t^{[in]}(S) = \min_r \{z_{rt} \mid z_{rt} \geq 0\}$ and the least negative inventory

$p_t^{[out]}(S) = \min_r \{-z_{rt} \mid z_{rt} < 0\}$. The suggested smoothing function $o_t(S)$ is

$$o_t(S) = \frac{1}{2N} \left(\frac{2p_t^{[in]}(S)}{p_t^{[in]}(S) + p_t^{[out]}(S)} - 1 \right). \quad (19)$$

It is proven in (Hendrix and Olieman, 2008), that $\hat{a}_t(S) + o_t(S)$ is continuous in the interesting values of S , as illustrated in Figure 4.1. Moreover, the function $\hat{a}_t(S) + o_t(S)$ deviates at most $\frac{1}{2N}$ from the reached service level $\hat{a}_t(S)$. This deviation is much smaller than the possible estimation error. Using $\hat{a}_t(S) + o_t(S)$ defines the problem $NLP(S, Y)$ where constraint (13) in MC-MILP is replaced by




Fig. 4.1 Illustration of SMC from (Hendrix and Olieman, 2008), where the estimated probability on the y-axis depends on varying one parameter on the x-axis

$$\hat{a}_t(S) + o_t(S) \geq \alpha \quad t = 1, \dots, T \quad (20)$$

as a smooth optimisation problem that in principle can be solved by a nonlinear optimisation routine. Notice again that only values S_t have to be determined for $Y_t = 1$ and $\exists i = 1, \dots, M-1, Y_{t-i} = 1$. For the chance constraints, one only has to focus on the last period of the replenishment cycle $t + R - 1$; the demand in between will have a higher probability to be fulfilled. As starting point to calculate values for the variables S_t in the nonlinear optimisation the values \hat{S}_{Rt} can be used. Algorithm 4.1 provides a list of order timing Y^* and order-up-to levels S^* that fulfil the chance constraints arbitrarily close if the number of samples N increases. One can use a lower bound on cost to decide that Y cannot be optimal. A lower bound $LBC(Y)$ on the cost contains the necessary minimum procurement cost $\Sigma(kY_t + c \cdot E(d_t))$. Moreover, the expected inventory at the beginning of a period where no order takes place is at least \hat{S}_{1t} and the corresponding inventory cost can be added to the lower bound $LBC(Y)$. In an enumeration of Y , if $LBC(Y)$ is greater than the best feasible objective value C^U found so far, Y cannot be the optimal timing.

Algorithm 4.1 *YSsmooth* in: samples d_{tr} , cost data, α , \hat{S}_{Rt} , out: Y^*, S^*

Set the best function value $C^U := \infty$

Generate the set of feasible order timing Y

for all Y

if for the lower bound on cost $LBC(Y) < C^U$

 solve $NLPS(Y)$ using \hat{S}_{Rt} values $\rightarrow S$ and cost C

if $C < C^U$

 save the best found values $C^U := C, S^* := S, Y^* = Y$

4.6 YQ(X) policy: sample based algorithm

We consider a static-dynamic uncertainty YQ(X) policy where the decision maker is provided an order timing vector Y , i.e. $Y_t = 0 \Rightarrow Q_t = 0$. The order quantity depends on the age distribution X of the items in stock. Practically, this is more complex than an order-up-to level for the decision maker as an information system advising on the order quantity given the actual age composition of the inventory is required. The YQ(X) policy is wider than the YS policy, so it should provide lower expected cost than the case where the age-distribution is not taken into account. However, it also suffers from the conditional way of considering the chance constraint, i.e. given the age distribution it can only focus on fulfilling the chance constraint in the future, no matter how likely the current age distribution X is, see (Rossi et al., 2008).

The characteristics of a solution found in Section 4.4 also hold for the YQ(X) policy. For the order moments fulfilling the conditions of Lemmas 2 and 3, the order quantities are provided by using the pre-calculated basic order-up-to levels \hat{S}_{Rt} .

For the other order moments, Lemma 4 tells us what to do when there is a negative inventory X . For a positive inventory X , the sample-based estimation of Section 4.4.3 can be used. At an order moment, i.e. $Y_t = 1$ where the inventory position is positive and $R > 1$, the order quantity may be larger than the basic order-up-to level

$$Q_t(X) \geq \left(\hat{S}_{Rt} - \sum_{b=1}^{M-1} X_{bt} \right)^+ \quad (21)$$

due to the occurrence of expected waste during the replenishment cycle. To compute the optimal order quantity for this case, we have to investigate the total inventory at the end of the replenishment cycle as function of the starting inventory X , the order quantity Q and the demand d_b, \dots, d_{t+R-1} during the replenishment cycle.

Definition 2. The function $Z: \mathbb{R} \times \mathbb{R}^M \times \mathbb{R}^R \rightarrow \mathbb{R}$ is defined as the transformation

$z = Z(Q, X, d)$ giving the total inventory $z = \sum_{b=1}^M I_{b,t+R-1}$ following the dynamics (4), (5) with starting inventory X , order quantity Q and demand vector (d_1, \dots, d_{t+R-1}) .

This definition facilitates writing the order quantity we are looking for as the minimum value Q_t for which the chance constraint holds; this is the value of Q_t for which $P(Z(Q_t, X_t, d_{t+R-1}) \geq 0) = \alpha$. The following property of function Z is useful.

Lemma 5. Let function Z be defined by Definition 2, $R \leq M$, values for Q, X, d given. Let $Z(Q, X, d) = z$, then $\forall q \in \mathbb{R}, Z(Q + q, X, d) = z + q$.

Proof. Due to $R \leq M$ following the equations (4), (5), nothing of the order quantity Q will be wasted (outdated). An additional amount q will be added to the total end inventory z . \square

Theorem 1. Let function Z be defined by Definition 2, $R \leq M$, starting inventory X given, $z = Z(0, X, d_b, \dots, d_{t+R-1})$ with cdf Γ . The optimal order quantity in period t is $Q_t = (-\Gamma^{-1}(1 - \alpha))^+$.

Proof. If $\Gamma^{-1}(1 - \alpha) > 0$, the current stock is enough to fulfil demand with probability α : $P(z \leq 0) < (1 - \alpha) \rightarrow P(z \geq 0) > \alpha$. So in that case, $Q_t = 0$ is optimal. For a value $Q_t = -\Gamma^{-1}(1 - \alpha) \geq 0$, we have $P(Z(0, X, d_b, \dots, d_{t+R-1}) \leq -Q_t) = 1 - \alpha$. This implies $P(Z(0, X, d_b, \dots, d_{t+R-1}) + Q_t \geq 0) = \alpha$. Using Lemma 5, this translates to $P(Z(Q_t, X, d_b, \dots, d_{t+R-1}) \geq 0) = \alpha$. So the order quantity Q_t is the minimum value for which the end inventory has a probability of α to be positive. Therefore it is the optimal value. \square

Lemmas 2 and 3 discussed the cases where $Q_t = \hat{S}_{Rt} - \sum_{b=1}^{M-1} X_{bt}$ is the optimal solution. A possible deviation from this value in other cases is due to the waste that can occur during the replenishment cycle. Taking this value as benchmark provides a corollary which follows directly from Theorem 1 and Lemma 5.

Corollary 1. Let function Z be defined by Definition 2, $R \leq M$, starting inventory X given, $z = Z(\hat{S}_{Rt} - \sum_{b=1}^{M-1} X_{bt}, X, d_t, \dots, d_{t+R-1})$ with cdf Γ . The optimal order quantity is $Q_t = \left(\hat{S}_{Rt} - \sum_{b=1}^{M-1} X_{bt} - \Gamma^{-1}(1 - \alpha) \right)^+$.

In other words, Lemmas 2 and 3 discuss cases where the choice $Q_t = \hat{S}_{Rt} - \sum_{b=1}^{M-1} X_{bt}$ gives $\Gamma^{-1}(1 - \alpha) = 0$. Notice, this is also the case if the starting inventory is non-positive, $X_{1t} \leq 0$, as no waste can be generated. In other cases, waste can be generated and $\Gamma^{-1}(1 - \alpha) < 0$. No analytical expression is available to evaluate its value. To estimate the quantile $\Gamma^{-1}(1 - \alpha)$, Monte Carlo simulation can be used as discussed in Section 4.4.3. Let D be an $N \times T$ matrix with samples d_{tr} . For a starting inventory X , giving the order quantity $Q_t = \hat{S}_{Rt} - \sum_{b=1}^{M-1} X_{bt}$, one can evaluate $z_r = Z(Q_t, X, d_{tr}, \dots, d_{t+R-1,r})$ being the total inventory of sample r at the end of the cycle. The adjusting amount $-\Gamma^{-1}(1 - \alpha)$ is estimated by

$$A_t(X) = (-\text{quantile}(\{z_r, r = 1, \dots, N\}, 1 - \alpha))^+, \quad (22)$$

where quantile $(\{\}, \alpha)$ is the α sample quantile of set $\{\}$.

The order quantity for any starting inventory according to Corollary 1, can be approximated by

$$Q_t(X) = \hat{S}_{Rt} - \sum_{b=1}^{M-1} X_{bt} + A_t(X). \quad (23)$$

This way of approaching the chance constraint is slightly stricter than the original service level constraints. It forces an α probability on positive inventory from any starting inventory

X , no matter what is the probability of being in this state. This is also called a conditional service level constraint, see (Rossi et al., 2008).

The order quantities for the $YQ(X)$ policy are now defined either by the theoretical results, or by the sample-based estimation in (22) and (23). The next question is to generate the best advice for the order timing Y . Algorithm 4.2 enumerates the possible timing vectors. Here one can make use of Lemma 1 and leave out those with too large periods between two orders. For each vector Y , the average cost is evaluated for a large simulation run that uses different random numbers than the ones in matrix D that are used to determine the order quantities by (22) and (23).

Algorithm 4.2 YQ in: samples d_{bt} , cost data, α , \hat{S}_{Rt} , out: Y^*

Set the best function value $C^U := \infty$

Generate the set of feasible order timing Y

for all Y

 if for the lower bound on cost $LBc(Y) < C^U$

 Determine C by simulating N sample paths

 During the simulation

 if $Y_t = 1$

 if starting inventory X not positive or $R_t = 1$ take $Q_t = \hat{S}_{Rt} - \sum_{b=1}^{M-1} X_{bt}$

 else simulate the replenishment cycle with N paths from X

 Determine the order quantity Q_t from (23)

 if $C < C^U$

 save the best found values $C^U := C$, $Y^* = Y$

The $YQ(X)$ policy has the advantage that it takes the age-distribution into account. However, for the decision maker the required use of tables and possibly interpolation is more hassle than using a simple order-up-to strategy with a list of order-up-to levels of the YS policy. The question is in which cases the policy $YQ(X)$ performs significantly better than the YS policy. The theoretical results already showed that in fact the $YQ(X)$ policy works with basic order-up-to levels \hat{S}_{Rt} for many cases. If costs and demand data are such that we order each period, or alternatively order every M periods, then in fact the $YQ(X)$ policy works with order-up-to levels all the time according to Lemmas 2 and 3.

4.7 Flexible $Q(X)$ policy

We consider a more flexible $Q(X)$ policy according to the dynamic uncertainty strategy. In this policy at the beginning of every period t , the order quantity is determined, based on the age-distribution of the available inventory. We obtain this policy by Stochastic Dynamic Programming (SDP), an appropriate technique to solve the SP model, as it is clearly separable in t . The SDP approach to solve the SP model is earlier described in (Hendrix et al., 2012).

The state values are given by X , the transition is provided by equations (4) and (5), so we have transition function Φ :

$$I_t = \Phi(X_t, Q_t, d_t) \quad t = 1, \dots, T \quad (24)$$

The chance constraints can be written as

$$Q_t \geq \left(\Gamma_t^{-1}(\alpha) - \sum_{b=1}^{M-1} X_{bt} \right)^+ \quad t = 1, \dots, T \quad (25)$$

The waste I_{Mt} is a function of the inventory at the beginning of the period and the demand: $I_{Mt} = f(X_t, d_t)$. We can write the expected contribution to the objective function in period t as function of state X_t and decision Q_t :

$$EC(X_t, Q_t) = g(Q_t) + E\{wf(X_t, d_t) + h1^T \Phi(X_t, Q_t, d_t)\} \quad (26)$$

where 1 is the all-ones vector. The SDP objective function can be written down via the Bellmann equation using a value function V :

$$V_t(X) = \min_Q (EC(X, Q) + E[V_{t+1}(\Phi(X, Q, d_t))]) \quad (27)$$

subject to fulfilling (25). The argmin of (27) represents the optimal policy $Q(X)$.

To implement this policy, at the beginning of every period, the decision maker needs an information system with the optimal strategy in tables, advising on the order quantity given the actual age composition of the inventory.

4.8 Comparison of policies

To compare the different policies, we use the same erratic demand pattern and the same design of experiments that is used in Chapter 2 (Pauls-Worm et al., 2014). The demand pattern is depicted in Table 4.2. In 81 experiments, we systematically vary fixed setup cost $k = 1500, 500, 2000$, disposal cost $w = -0.5, 0, 0.5$, α -service level = 90%, 95%, 98%, $CV = 0.1, 0.25, 0.333$. The other values are kept constant: the product has a shelf life of $M = 3$, unit production cost $c = 2$ and unit holding cost $h = 0.5$. Negative disposal cost means the product has a salvage value, which is usually much less than the unit production cost c , zero disposal cost means that only the unit production cost is lost in case of waste, and positive disposal cost means that there is a cost to discard the wasted items.

We compare a YS_{MILP} policy with parameters generated by MILP (Pauls-Worm et al., 2014) (Chapter 2), a $YS_{SMC-MINLP}$ policy with parameters generated by the smoothed Monte Carlo MINLP approach described in Section 4.5, a $YQ(X)$ policy according to the sample based algorithm presented in Section 4.6, and the flexible $Q(X)$ policy generated by SDP as discussed in Section 4.7. The inventory system is simulated for all policies using the same (pseudo) random number series of 10,000 runs and compared on expected total costs with the expected total costs of the $YS_{SMC-MINLP}$ policy set to 100 as reported in Table 4.3. The YS_{MILP} policy does not always meet the required service level, especially at the end of the horizon. The developed $YS_{SMC-MINLP}$ policy always meets the required service levels, but requires a much longer computation time.

Table 4.2 Erratic demand pattern

t	1	2	3	4	5	6	7	8	9	10	11	12
Data $E(d_t)$	800	950	200	900	800	150	650	800	900	300	150	600

In Table 4.3 we can observe that in most scenarios the expected total costs are very close. The YS_{MILP} policy provides often the lowest cost, but does not always meet the service level requirements. However, this is mostly in the last period $T = 12$ due to end-of-horizon effects. The $YS_{MC-MINLP}$ policy is in general the policy that meets the service level best at lowest cost. The shaded cells show the scenarios where a more complex policy is preferred. In case of low setup cost ($k = 500$), one orders almost every period the expected total costs of the different order policies are almost equal. Only with increasing uncertainty and higher service levels the $Q(X)_{SDP}$ might significantly save costs. However, the dynamic uncertainty strategy $Q(X)_{SDP}$ policy may raise the expense on planning, which is not part of this model. In case of high setup cost ($k = 2000$), the static-dynamic uncertainty YS and $YQ(X)$ policies have mostly the same production plans based on the derived basic order-up-to levels and therefore, the costs show no significant differences. The $Q(X)_{SDP}$ policy has in most cases significantly higher costs. This is due to the overachievement of the service level. SDP meets a conditional service level requirement as is shown in (Rossi, 2013a) and (Pauls-Worm and Hendrix, 2015) (Chapter 3) and is allowed to order every period, which is very costly in case of high setup cost.

The $YQ(X)$ policy also meets a conditional service level requirement, but here the production moments are fixed, which results in a less “nervous” system (Tunc et al., 2013). Due to this behaviour, in scenarios 80 and 81 the $YQ(X)$ policy has lower costs, comparable to the costs of $Q(X)_{SDP}$. This means the $YQ(X)$ policy is preferred as it fixes the order timing. As also shown in Chapter 2 (Pauls-Worm et al., 2014), an intermediate level of setup cost is the most interesting situation. The replenishment cycles are varying in length, resulting in waste during the replenishment cycles. Finding the optimal order timing is more difficult as discussed in Section 4.4, and considering the age-distribution of the items in stock becomes more important. This is confirmed by the costs of the different order policies, in the scenarios with higher uncertainty and higher service levels (16 – 18, 22 – 27). The $YQ(X)$ policy, taking the age-distribution of the items in stock into account, gives clearly lower costs, as expected.

Table 4.3 Simulated costs results of four order policies relative to $YS_{SMC-MINLP}$ costs = 100

$k = 1500$	CV = 0.10				CV = 0.25				CV = 0.33			
	Scenario	YS_{MILP}	$YQ(X)$	$Q(X)_{SDP}$	Scenario	YS_{MILP}	$YQ(X)$	$Q(X)_{SDP}$	Scenario	YS_{MILP}	$YQ(X)$	$Q(X)_{SDP}$
Service level 90%												
$w = -0.5$	1	100.03	100.12	102.56	10	99.86	100.28	103.01	19	99.13	100.34	102.83
$w = 0$	2	100.04	100.14	102.12	11	99.51	99.86	102.18	20	98.80	99.57	101.82
$w = 0.5$	3	99.74	100.16	101.59	12	99.31	99.99	101.66	21	98.49	99.35	101.24
Service level 95%												
$w = -0.5$	4	100.02	100.13	101.61	13	99.24	99.68	101.22	22	98.76	97.20	100.77
$w = 0$	5	100.02	100.15	100.96	14	99.12	99.90	100.82	23	98.23	98.88	99.76
$w = 0.5$	6	99.40	99.83	100.27	15	98.81	99.39	100.41	24	97.13	97.84	98.40
Service level 98%												
$w = -0.5$	7	100.04	100.18	100.61	16	98.79	96.98	99.51	25	98.70	98.64	98.66
$w = 0$	8	99.93	100.21	99.72	17	98.23	98.89	99.11	26	98.51	98.41	98.02
$w = 0.5$	9	99.55	99.80	99.03	18	97.99	98.72	98.72	27	98.44	98.29	97.43

$k = 500$	CV = 0.10				CV = 0.25				CV = 0.33			
	Scenario	YS_{MILP}	$YQ(X)$	$Q(X)_{SDP}$	Scenario	YS_{MILP}	$YQ(X)$	$Q(X)_{SDP}$	Scenario	YS_{MILP}	$YQ(X)$	$Q(X)_{SDP}$
Service level 90%												
$w = -0.5$	28	99.95	99.99	100.95	37	99.66	99.91	100.17	46	99.33	99.88	99.28
$w = 0$	29	99.55	99.99	100.76	38	99.65	99.90	100.15	47	99.32	99.86	99.20
$w = 0.5$	30	99.63	99.98	100.90	39	99.65	99.89	100.11	48	99.30	99.84	99.10
Service level 95%												
$w = -0.5$	31	99.56	100.00	100.30	40	99.56	99.94	99.27	49	99.18	99.88	98.39
$w = 0$	32	99.67	99.98	100.46	41	99.56	99.94	99.23	50	99.17	99.87	98.27
$w = 0.5$	33	99.77	99.98	100.59	42	99.55	99.92	99.18	51	99.16	99.86	98.15
Service level 98%												
$w = -0.5$	34	99.96	100.02	99.89	43	99.50	100.03	98.43	52	99.14	100.04	97.60
$w = 0$	35	99.96	100.01	100.06	44	99.49	100.02	98.36	53	99.13	100.03	97.44
$w = 0.5$	36	99.96	100.01	100.21	45	99.48	100.01	98.29	54	99.13	100.02	97.29

$k = 2000$	CV = 0.10				CV = 0.25				CV = 0.33			
	Scenario	YS_{MILP}	$YQ(X)$	$Q(X)_{SDP}$	Scenario	YS_{MILP}	$YQ(X)$	$Q(X)_{SDP}$	Scenario	YS_{MILP}	$YQ(X)$	$Q(X)_{SDP}$
Service level 90%												
$w = -0.5$	55	100.03	100.11	103.34	64	100.02	100.26	105.04	73	100.01	100.32	105.15
$w = 0$	56	100.03	100.13	103.49	65	100.03	100.30	104.16	74	99.92	100.37	103.93
$w = 0.5$	57	100.04	100.15	103.45	66	100.21	100.52	103.34	75	99.09	100.42	102.55
Service level 95%												
$w = -0.5$	58	100.02	100.12	102.46	67	100.02	100.27	103.41	76	100.01	100.34	103.14
$w = 0$	59	100.02	100.14	102.46	68	100.02	100.31	102.15	77	99.41	99.95	101.54
$w = 0.5$	60	100.02	100.16	102.23	69	99.19	100.35	100.83	78	98.76	99.10	100.05
Service level 98%												
$w = -0.5$	61	100.04	100.16	101.66	70	99.97	100.32	101.78	79	100.02	100.44	101.22
$w = 0$	62	100.04	100.20	101.50	71	99.31	100.37	100.08	80	99.31	98.71	98.78
$w = 0.5$	63	100.05	100.23	100.86	72	98.68	99.11	99.04	81	98.55	98.42	97.99

4.9 Conclusion

We developed new sample based order policies (control rules) for a SP inventory control problem in case of non-stationary demand and compared the new policies with existing policies. The inventory control can be handled according to a static-dynamic uncertainty strategy, where we distinguish a YS policy and a $YQ(X)$ policy. In inventory literature of perishable products this strategy is little exposed. In this paper we first discussed characteristics of a solution to these policies. Next we presented a new computational

method based on the Smoothed Monte Carlo method with sampled demand, called the $YS_{SMC-MINLP}$ policy and a sample based method to calculate values for the $YQ(X)$ policy. These policies were compared to an YS_{MILP} policy and a more flexible $Q(X)$ policy generated by SDP, according to a dynamic uncertainty strategy. The experimental evaluation comprises 81 scenarios with the same erratic demand pattern, but with varying setup cost, service level, cost of waste and uncertainty measured in the Coefficient of Variation.

In most scenarios the expected total costs of the policies are very close, and a YS policy gives a cost-efficient and easy to implement solution. From a production planning perspective, the static-dynamic uncertainty strategy is the most convenient strategy to follow. In situations of relatively low setup cost, where every period is produced, or high setup cost, where every M periods is produced, MILP generates appropriate parameter values. In situations of intermediate setup cost, where the replenishment cycles are highly varying, the optimised SMC-MINLP parameters might be more suitable, although more calculation time is needed. It is up to management to decide whether this is worthwhile. When also the required service level is high or the CV is 0.25 or more, a $YQ(X)$ policy, where the age-distribution of the inventory is taken into account, might have less inventory costs. However, for the decision maker the required use of tables and possibly interpolation is more hassle than using a simple order-up-to strategy with a list of order-up-to levels of the YS policy. The dynamic uncertainty strategy is only appropriate in situations with relatively low setup cost, a CV of 0.33 and a high service level of 98%. The same implementation drawback as for the $YQ(X)$ policy applies.

Chapter 5

YS policy in retail

Based on:

Pauls-Worm, K.G.J., Hendrix, E.M.T., 2016. *On retail order policies for a perishable product*, Submitted to an international journal.

Abstract

We study retail order policies for a practical situation where stores may be able to order every day, or on fixed days of the week. Focus is on a perishable product which has a remaining shelf life on delivery at the store of three days and a weekly (seasonal) demand pattern which is stationary over the weeks. The stores have a target α -service level, i.e. the inventory should cover demand in more than α % of the days. The customer demand may be First In First Out (FIFO) or a Last In First Out (LIFO) – FIFO combination. A Stochastic Programming (SP) model is presented of the situation in the retailer practice. Four different policies to determine the order quantity are studied. The base is a YS order policy where the reorder days Y are fixed and order-up-to levels S are used, with parameter values generated by an MILP approximation of the SP model. Numerical experiments compare the effectiveness of the policies to determine the order quantities with respect to costs and reached service levels. The first policy determines the order quantity considering the total available inventory. This gives lowest cost solutions, which are not always feasible. From the three other studied policies, the policy where a fraction of the total available inventory is considered when determining the order quantity, performs best. One can always find a fraction that gives feasible solutions.

5.1 Introduction

This study is based on the practical situation at a Dutch retail organisation that has stores of various sizes, each with potentially different optimal order policies. The retail organisation wants to improve the ordering of fresh food products. For this case we study the inventory system of a perishable product with a remaining shelf life on delivery at the store of three days. The product has a weekly (seasonal) demand pattern which is stationary over the weeks. Stores may be able to order every day, or on fixed days of the week.

Nahmias (1982), Goyal and Giri (2001), and Karaesmen et al. (2011) reviewed order policies for perishable products with a fixed lifetime. Almost all papers surveyed assume stationary demand. Nahmias (1975) and Fries (1975) note that generally an optimal order policy for perishable products with a fixed life time should consider the age-distribution of the products in stock. Van Donselaar et al. (2006) studied the characteristics of perishable products, their supply, sales and inventory control in supermarkets. They found a weekly pattern for sales in Dutch supermarkets, with higher sales on Friday and Saturday. At that time, store opening on Sunday was less common in the Netherlands. All studies we found about order policies for a fixed lifetime perishable product in a (potential) retail application e.g. (Broekmeulen and Van Donselaar, 2009), (Minner and Transchel, 2010), (Chao et al., 2015), (Haijema and Minner, 2015), assume a periodic review, where a reorder decision can be made in every period. Withdrawal of product by consumers is either FIFO (First In First Out), LIFO (Last In First Out) or a combination of FIFO and LIFO. Broekmeulen and Van Donselaar (2009) study a periodic review policy for a retail situation which takes the age of the inventory in the system into account. They assume a positive lead time, a weekly (seasonal) demand pattern which is stationary over the weeks. The safety stock is constant. The review period is R days, meaning that every R ($R = 1$ or 2) days, the store has the opportunity to order one or more batches. A batch is a set of items with the same remaining shelf life. Broekmeulen and Van Donselaar (2009) introduce a new policy, called EWA, in which they determine the order quantity by correcting the inventory position for the estimated outdating. Outdating is estimated by calculating the withdrawal from stock by the expected demand, the remaining batches and the outdating. Withdrawal is either FIFO or LIFO. Broekmeulen and Van Donselaar (2009) evaluate service levels, but they are not part of the model formulation. Minner and Transchel (2010) present a service level approach, with α -service level constraints and fill rate constraints, for a retail situation where they determine dynamic order quantities for perishable products with a fixed maximum shelf life, positive lead time and FIFO or LIFO withdrawal. In their periodic review system, an order can be placed every day. They assume a weekly (seasonal) demand pattern which is stationary over the weeks. Chao et al. (2015) introduce approximation algorithms for perishable inventory systems with a worst-case cost performance guarantee, for non-stationary correlated demand processes and for independent and stochastically non-decreasing demand processes. They assume FIFO withdrawal of the items. For the latter case, they propose a dual-balancing policy where the order quantity is determined by balancing the expected discounted marginal holding and outdating costs with the expected marginal backlogging cost. This turns out to be an RS_t policy. Haijema and Minner (2015) compare in a periodic review, hybrid base-stock and constant order policies for a perishable

product with a stationary demand in a retail situation. They assume combined FIFO and LIFO withdrawal. They show that a hybrid policy performs better than a traditional base-stock policy.

To our knowledge, all papers study policies according to the dynamic uncertainty strategy (Bookbinder and Tan, 1988) where the order quantity is decided at the beginning of every day or every two days. However, in practice, at smaller stores, or for low demand products, it might not be optimal or desirable to order fresh products every day. Every day order and delivery may increase costs, and may be undesirable from an environmental perspective, such as CO₂ emissions, noise and hindrance of traffic. Moreover, every day delivery increases the number of batches with different remaining shelf life in stock, which may complicate the inventory management. Store managers try to get the customers buy FIFO by putting the older items in front on the shelf. However, customers who are determined to buy the freshest items (LIFO) will pick from the back of the shelf, thus messing up the inventory on the shelf. Store clerks will regularly have to reorganise the inventory on shelf. Ordering every two days causes different reorder schedules per week. For planning purposes of the distribution centre and the store, reordering on fixed days of the week is desirable and common in practice. The question is on which days to order and how much. Due to the weekly demand pattern and the perishability of the product, the length of the replenishment cycles in a week may vary in length. In retail, stock-out is easily established, whilst the amount of lost sales is unknown. Hence, an α -service level is applied, defined as the probability of no stock-out at the end of the day. The stores in our study have a target α -service level of 90%. Very often items of different ages are in store. Stores can influence the picking of the items by the customers, so we study FIFO withdrawal. However, some customers pick the items in front, and some search for the freshest items. Therefore we also study combined LIFO – FIFO withdrawal. The aim of the study is to find a suitable reorder plan according to a static-dynamic uncertainty strategy (Bookbinder and Tan, 1988) where the reorder days are fixed. A YS policy where $Y_t = 1$ denotes a reorder day and S_t the corresponding order-up-to level, is an easy to implement policy.

Pauls-Worm et al. (2014) (Chapter 2) present an MILP model which generates parameter values for a YS order policy for perishable items under an α -service level constraint with zero lead time as a solution to an SP model of the problem. In this chapter we adapt this SP model to the weekly demand pattern situation and derive an MILP approximation for a lead time of one day and a combined LIFO and FIFO demand by the customers, to model the situation in the retailer practice. We study the practical application of these models simulating a rolling horizon. The MILP model determines the replenishment cycle lengths with the corresponding order-up-to levels. These parameters are input for the rolling horizon approach. The research question is: are the MILP generated parameter values suitable in case of a lead time of one period and FIFO or combined LIFO – FIFO withdrawal in a rolling horizon application?

In Section 5.2 we describe the model assumptions, followed by the SP model in Section 5.3. In Section 5.4 four order policies are introduced, for which Section 5.5 presents the MILP

numerical experiments with a simulation of the four policies. Section 5.7 summarises the findings of the investigation.

5.2 Model assumptions

In this study a period t in the model is a day at the store, from opening until closing time. In the retail practice of perishable products, mostly the order Q_t of today is delivered tomorrow, so we use a lead time $L = 1$. An order is placed at a fixed setup cost k . In the model the sequence of events is as follows:

- Store opening
- Delivery Q_{t-1}
- Order Q_t
- Demand during the day, Poisson distributed
- Update ages of items and dispose of wasted items at store closing

On the moment of the order decision, the previous order has arrived, so there is no outstanding order. The order quantity is based on the on-hand inventory and the forecasted demand for the upcoming periods. At the end of the day, the inventory level I_{bt} of items of all ages b is checked. Items that are delivered at day t in quantity Q_{t-1} , have age $b = 1$ at the end of day t . Items with an age that reaches the shelf life $b = M$, are waste and removed from the shelf at the end of the day, at a disposal cost of w per unit. A holding cost h is incurred for items that are in the store overnight. The time horizon T is 7 days, where $t = 1$ is Monday. Retail stores tend to get fixed days on which they are allowed to reorder at the distribution centre. Therefore we assume the expected inventory levels at the beginning of the week to be equal to the expected inventory levels at the end of the week, in order to obtain a fixed weekly reorder plan with order-up-to levels for each order day. The value of the order-up-to level S_t has to be such that the demand of day t and $t + 1$ can be met, if one orders the next day again. In case of a shelf life of M , one has to order at least every M days, using an order-up-to level S_t to meet demand of d_b, \dots, d_{t+M} , so for $M + 1$ days. The demand is independently Poisson distributed with expectation $\mu_t \in \mathbb{R}$ for day t with a value obtained from a practical case. In the following section, we present a general SP model for this problem.

5.3 Stochastic Programming model

The reorder decision problem can be formulated as a stochastic programming model that minimises ordering and purchasing costs, holding cost and cost of waste.

$$\text{Min } E(TC) = \int \dots \int \sum_{d_T, t=1}^T \left(f(Q_t) + h \sum_{b=1}^{M-1} I_{bt} + wI_{Mt} \right) g_1(d_1) \dots g_T(d_T) dd_1 \dots dd_T \quad (1)$$

The expected total costs are modelled as a continuous function, where the function $g_t(d_t)$ can be any probability density function of the demand. Function $g_t(d_t)$ can also represent a discrete probability mass function, like a Poisson distribution.

The ordering and purchasing costs are given by a fixed setup cost k on ordering and a variable purchasing cost c per unit, i.e.

$$f(Q) = \begin{cases} k + cQ & \text{if } Q > 0 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

The inventory balance constraint for the total inventory of all ages is given by:

$$\sum_{b=1}^M I_{bt} = \sum_{b=1}^{M-1} I_{b,t-1} + Q_{t-L} - d_t \quad t = 1, \dots, T \quad (3)$$

Period t starts with the inventory levels at the end of period $t - 1$ of ages $b = 1, \dots, M - 1$, since items of age M are waste. The starting inventory is increased by the delivery Q_{t-L} minus the demand in period t , giving the end inventory.

$$P\left(\sum_{b=1}^M I_{bt} \geq 0\right) \geq \alpha \quad t = 1, \dots, T \quad (4)$$

The probability that the sum of the inventory levels at the end of period t of all ages is nonnegative should be at least α . This type of service level is known as α -service level. Let $(x)^+ = \max\{x, 0\}$, then the inventory dynamics in a FIFO situation, where demand is fulfilled first by the oldest items before the fresher items, can be described as follows.

$$I_{bt} = \left(I_{b-1,t-1} - \left(d_t - \sum_{j=b}^{M-1} I_{j,t-1} \right)^+ \right)^+ \quad t = 1, \dots, T; \quad b = 2, \dots, M \quad (5)$$

describe the levels of waste and the older items in stock.

$$I_{1t} = Q_{t-L} - \left(d_t - \sum_{b=1}^{M-1} I_{b,t-1} \right)^+ \quad t = 1, \dots, T \quad (6)$$

give the freshest items in stock.

Equations (7) and (8) model the inventory levels in case of only LIFO withdrawal, where demand is fulfilled first by the freshest items before the older items.

$$I_{1t} = (Q_{t-L} - d_t)^+ \quad t = 1, \dots, T \quad (7)$$

$$I_{bt} = \left(I_{b-1,t-1} - \left(d_t - Q_{t-L} - \sum_{j=1}^{b-2} I_{j,t-1} \right)^+ \right)^+ \quad t = 1, \dots, T; \quad b = 2, \dots, M \quad (8)$$

In case of only FIFO withdrawal of demand, there can only be a shortage of the freshest items, but in case of only LIFO withdrawal of demand there can be a shortage of the oldest items. When picking of demand is combined LIFO and FIFO, we assume that LIFO demand is met first. First a fraction λ of demand is realised for the freshest items (LIFO) with $0 \leq \lambda \leq 1$. After that, a fraction $1 - \lambda$ of demand is issued according to FIFO.

$$I_{b0} = I_{bT} \quad b = 1, \dots, M \quad (9)$$

$$I_{bt} \geq 0 \quad t = 1, \dots, T; \quad b = 1, \dots, M \quad (10)$$

$$Q_{1-L} = Q_{T+1-L} \quad (11)$$

$$Q_t \geq 0 \quad t = 1, \dots, T \quad (12)$$

The starting inventory level in period 1 of all ages equals the inventory level of all ages at the end of the week (9) and the inventory levels of all ages in all periods are nonnegative (10). In case demand exceeds the inventory, the excess demand is lost and the inventory levels are zero. The last order quantity Q_{T+1-L} of the time horizon equals the first delivery Q_{1-L} of the time horizon (11). The order quantity is nonnegative (12).

5.4 Order policies

In retail practice, ordering is on fixed days of the week. A possible policy is the so-called *YS* policy, where the reorder days are fixed and an order-up-to level is used. Finding suitable parameter values for a *YS* policy with MILP is interesting for use in practice, because it is a very fast way to determine a policy. In Section 5.5 we present a deterministic MILP model that generates parameter values for this policy. The stores have a target α -service level, therefore we want to find the value S_t for which the non-stockout probability is α . For the Poisson distributed demand this is a value $S \in \{0, 1, 2, 3, \dots\}$, such that

$$P(d \leq S) = e^{-\mu} \sum_{i=0}^S \frac{\mu^i}{i!} \geq \alpha \quad (13)$$

As can be found in standard textbooks like (Chopra and Meindl, 2010), in a periodic review system the order-up-to level consists of the demand during the lead time L and the replenishment cycle $r \leq M$ increased with the safety stock for the lead time and the replenishment cycle. The Poisson distribution allows us to calculate order-up-to levels for all possible replenishment cycle lengths with lead time by summing the expected demands and safety stocks. We call these order-up-to levels the basic order-up-to level $\hat{S}_{L+r,t}$ for ordering for r periods in period t as described in (Hendrix et al., 2015). The values for S can easily be derived using an Excel or Matlab search routine. The order quantity in terms of an order-up-to policy is determined by

$$Q_t = S_t - Q_{t-L} - \sum_{b=1}^{M-1} I_{b,t-1} \quad (14)$$

During a replenishment cycle two issues might occur. First, the described concept to include the lead time demand in the order-up-to level is based on a backlogging situation, which is not reflecting a retail stock-out situation. When at the start or during period t a stock-out occurs, the order quantity Q_t will be higher than necessary. Secondly, in a perishable inventory situation waste may occur during the replenishment cycle. The S_t value obtained from the MILP model is the basic order-up-to level $\hat{S}_{L+r,t}$ with added expected waste. However, in Chapter 2 (Pauls-Worm et al., 2014) explained that in the MILP model the expected waste is underestimated. The S_t value obtained from the MILP model may not be high enough. These issues provide a challenge to develop a practical good policy.

In case of lost sales during the lead time, one can determine the order quantity based on the following reasoning. To simplify notation we explicitly use the lead time $L = 1$, although the results can easily be extended. Let $F_{r,k}()$ be the cumulative distribution function of

demand $d_{t+1} + \dots + d_{t+r}$ during the replenishment cycle, then $\hat{Q}_{rt} = F_{rt}^{-1}(\alpha)$ gives the amount that should at least be in stock at the beginning of the next day. In an out-of-stock situation in period t , this quantity is exactly the amount to be ordered. Fig. 5.1 shows the time frame of the used symbols.




Fig 5.1 Overview of the time frame of the used symbols and possible age-distribution in stock for a replenishment cycle of $r = 3$ periods

Lemma 1. Let the starting inventory $\sum_{b=1}^{M-1} I_{b,t-1} + Q_{t-1} = 0$ and r be the length of the replenishment cycle. The optimal order quantity is $Q_t = \hat{Q}_{rt}$.

Proof. Minimisation of the costs leads to a value of Q_t as low as possible. To fulfil chance constraint (5), the order quantity should fulfil $Q_t \geq \hat{Q}_{rt}$. Minimising its value implies $Q_t = \hat{Q}_{rt}$.

Based on the above properties for S_t and Q_b we propose three alternative policies to calculate the order quantity based on the order-up-to level S_b and an alternative way to calculate Q_b . Let us notice, that in the practical case our study is based on, the age-distribution of the items in stock is not known. In the first policy, the total available inventory at the beginning of period t is considered to determine the order quantity.

$$Q_t = S_t - Q_{t-1} - \sum_{b=1}^{M-1} I_{b,t-1} \quad (15)$$

When at the start or during period t a stock-out occurs, the order quantity Q_t will be higher than necessary. A practical way to deal with the lost sales during lead time is:

$$\text{if } Q_{t-L} + \sum_{b=1}^{M-1} I_{b,t-1} - \mu_t \leq 0, \text{ then } Q_t = \hat{Q}_{rt} \quad (16)$$

extending the finding of Lemma 1. We call this policy $YS_{\Sigma t}$.

The order-up-to level S_t generated by the MILP model is corrected for the expected waste, but the expected waste is underestimated. To take care of this phenomenon, we determine the order quantity Q_t considering only the freshest items in stock, under the condition of expression (16).

$$Q_t = S_t - Q_{t-L} - \sum_{b=1}^{M-2} I_{b,t-1} \quad (17)$$

We call this policy YS_{Fresh} . This policy assumes that the store can identify the number of oldest items in stock. In the practical case that motivated our study, the checkout systems of the supermarket do not register the age of the items sold, so in the supermarket one does not know the age-distribution of the perishable items in stock, unless visual inspection takes place. Taking this into account, the third policy is to determine the order quantity Q_t by considering a fraction γ of the available inventory from previous periods, following the additional rule of expression (16).

$$Q_t = S_t - Q_{t-L} - \gamma \sum_{b=1}^{M-1} I_{b,t-1} . \quad (18)$$

The manager at the retail store can adapt the parameter γ depending on the reached service level. In our experiments, the fraction γ is determined such that for every fraction of LIFO demand a feasible solution is reached in most scenarios. We call this policy $YS_{\gamma \Sigma t}$.

Based on the optimal order quantity \hat{Q}_{rt} for the case of zero starting inventory, we define a fourth order policy. A positive stock situation at the end of period t provides two challenges

to find the best order quantity. First, one has to predict the stock situation $\sum_{b=1}^{M-1} I_{bt}$ at the

beginning of period $t+1$. Second, waste $I_{M,t+1}, \dots, I_{M,t+r}$ can occur during the replenishment cycle. Fig. 5.1 shows when the available inventory at the start of period t becomes waste during a replenishment cycle of three periods. Lemma 1 is an aid to express the optimal order quantity

$$Q_t = \hat{Q}_{rt} - E \sum_{b=1}^{M-1} I_{bt} + E \sum_{j=t+1}^{t+r} I_{Mj} . \quad (19)$$

However, both expectations are difficult to determine and even complicated in the SP model due to the mixed LIFO – FIFO customer behaviour. Moreover, the age-distribution of the items in stock, necessary to predict both expectations, may be unknown in the practical situation. One can determine in an exact way the expected part of the delivered quantity

Q_{t-1} which is still in stock at the end of the lead time, $E(I_{1t} | Q_{t-1}, \sum_{b=1}^{M-1} I_{b,t-1})$. A practical way

to deal with the order quantity is to ignore the rest of the inventory and take as order quantity

$$Q_t = \hat{Q}_{rt} - E(I_{1t} | Q_{t-1}, \sum_{b=1}^{M-1} I_{b,t-1}). \quad (20)$$

This quantity is easy to determine from the probability distribution and the LIFO – FIFO behaviour, but takes some calculation time. The decision when to order (Y_t) is taken from the MILP model. We call this policy $Y\hat{Q}_{rt}$.

5.5 MILP model

The parameter values of a YS policy are derived by an MILP approximation with more strict assumptions as the model described in Chapter 2 (Pauls-Worm et al., 2014), like $I_{b0} = I_{bT}$ and $Q_0 = Q_T$ in case $L = 1$, but extended with a lead time L and a combined LIFO – FIFO demand. To determine values for Y_b when to order, and values for the order-up-to level S_b , expected values for the demand μ_b , the inventory levels of all ages I_{bt} and the order quantity Q_t are used. The expected total costs are minimised.

$$\text{Min } E(TC) = \sum_{t=1}^T \left\{ kY_t + h \sum_{b=1}^{M-1} EI_{bt} + cEQ_t + wEI_{Mt} \right\} \quad (21)$$

$$EQ_t \leq \mathcal{M}_t Y_t \quad t = 1, \dots, T \quad (22)$$

\mathcal{M}_t is a sufficiently large number, for instance $\mathcal{M}_t = \hat{S}_{L+M,t}$. This upper bound is high enough, as the order quantity in the deterministic MILP model will not exceed the necessary quantity to cover M periods.

The model formulation of the α -service level implies a minimal level of safety stock at the end of the replenishment cycle. The safety stock can be determined by taking the basic order-up-to level and subtracting the demand during the lead time and the replenishment cycle. In this approach the MILP model calculates the required order-up-to level, which is the basic order-up-to level $\hat{S}_{L+r,t}$ or equivalently expressed as $\hat{S}_{t-L-j+1,t}$, if necessary increased with the expected waste during the replenishment cycle. The binary variable $Z_{jt}=1$ indicates that the most recent order prior to period t is to meet demand for j periods. The inventory level at the end of period t has to be at least the safety stock to meet demand of the last j period(s) and the lead time. Let $G_t(\alpha)$ be the cumulative distribution function of demand.

$$\sum_{b=1}^M EI_{bt} \geq \sum_{j=1}^M \left(G_{t-L-j+1,t}^{-1}(\alpha) - \sum_{n=t-L-j+1}^t \mu_n \right) \cdot Z_{jt} \quad t = 1, \dots, T \quad (23)$$

There is only 1 most recent order prior to period t , to meet demand for at most M periods.

$$\sum_{j=1}^M Z_{jt} = 1 \quad t = 1, \dots, T \quad (24)$$

If the most recent order prior to period t is to meet demand for 1 period, there is an order in period $t-L$,

$$Z_{1t} = Y_{t-L} \quad t = 1, \dots, T \quad (25)$$

If the most recent order prior to period t is to meet demand for j periods, there are no orders between period $t - L - j + 1$ and $t - L + 1$.

$$Z_{jt} \geq Y_{t-L-j+1} - \sum_{n=t-L-j+2}^{t-L} Y_n \quad t = 1, \dots, T; \quad j = 2, \dots, M \quad (26)$$

We assume that the time horizon is repeated, so when $t = 1$, then $t - 1 = T$. The expected order quantity equals the order-up-to level minus the delivered quantity and the inventory on hand at the end of the previous period.

$$EQ_t = S_t - EQ_{t-L} - \sum_{b=1}^{M-1} EI_{b,t-1} \quad t = 1, \dots, T \quad (27)$$

The items $EI_{M,t-1}$ are left out, since they became waste and cannot be used in period t . To determine the order-up-to levels, it is necessary to keep track of the expected age-distribution of the inventory. When withdrawal of demand is combined LIFO and FIFO, we assume that LIFO demand is met first. The fraction of demand for the freshest items (LIFO) is denoted by I with $0 \leq I \leq 1$. The fraction of demand that is issued FIFO is denoted by $1 - I$. The auxiliary variable EI_{bt} is the expected inventory of age b after LIFO demand. Fresh items from the delivered order quantity have age 0. Let the auxiliary variables EXI_{bt} and EX_{bt} denote the residual demand for items of age b with $b = 1, \dots, M - 1$ in period t , for respectively LIFO and FIFO demand. If EXI_{bt} or EX_{bt} has a positive value, then older respectively fresher inventory is used to fulfil demand. At the end of the day, after FIFO demand, the ages are updated.

$$EQ_{t-L} - I \cdot \mu_t = EI_{0t} - EXI_{0t} \quad t = 1, \dots, T \quad (28)$$

$$EI_{b,t-1} - EXI_{b-1,t} = EI_{bt} - EXI_{bt} \quad t = 1, \dots, T; \quad b = 1, \dots, M - 1 \quad (29)$$

FIFO demand:

$$EI_{M-1,t} - (1 - I) \cdot \mu_t - EXI_{2t} = EI_{Mt} - EX_{M-1,t} \quad t = 1, \dots, T \quad (30)$$

$$EI_{bt} - EX_{b+1,t} = EI_{b+1,t} - EX_{bt} \quad t = 1, \dots, T; \quad b = 0, \dots, M - 2 \quad (31)$$

The right-hand-sides of equations (28) to (31) can each contain at most one variable with a positive value. The other variable needs to have a value of 0. Equations (32) to (35) impose that, using the binary variables BXI_{bt} and BX_{bt} .

$$\mathcal{M} \cdot BXI_{bt} \geq EXI_{bt} \quad t = 1, \dots, T; \quad b = 0, \dots, M - 1 \quad (32)$$

$$\mathcal{M} \cdot (1 - BXI_{bt}) \geq EI_{bt} \quad t = 1, \dots, T; \quad b = 0, \dots, M - 1 \quad (33)$$

$$\mathcal{M} \cdot BX_{bt} \geq EX_{bt} \quad t = 1, \dots, T; \quad b = 1, \dots, M - 1 \quad (34)$$

$$\mathcal{M} \cdot (1 - BX_{bt}) \geq EI_{b+1,t} \quad t = 1, \dots, T; \quad b = 1, \dots, M - 1 \quad (35)$$

The starting inventory level in period 1 of all ages equals the end inventory level of all ages in period T (36) and the delivered order in period 1 is ordered in period $T + 1 - L$ (37).

$$EI_{b0} = EI_{bT} \quad b = 1, \dots, M \quad (36)$$

$$EQ_{1-L} = EQ_{T+1-L} \quad (37)$$

$$EI_{bt}, EQ_t, S_t \geq 0 \quad t = 1, \dots, T; \quad b = 1, \dots, M \quad (38)$$

$$EI_{bt}, EXI_{bt}, EX_{bt} \geq 0 \quad t = 1, \dots, T; \quad b = 1, \dots, M - 1 \quad (39)$$

$$Y_t, Z_{tj} \in \{0, 1\} \quad t = 1, \dots, T; \quad j = 1, \dots, M \quad (40)$$

$$BXI_{bt}, BX_{bt} \in \{0, 1\} \quad t = 1, \dots, T; \quad b = 1, \dots, M - 1 \quad (41)$$

Equations (38) to (41) are definition constraints.

5.6 Numerical experiments

The design of experiments is described in Section 5.6.1, followed by the results for 27 scenarios in Section 5.6.2. Finally, in Section 5.6.3 one scenario is discussed in more detail to get insight in the dynamics of the inventory system.

5.6.1 Design of experiments

We apply the solution of the MILP model in a rolling horizon simulation of 10,000 weeks. We consider an expected demand varying with the day of the week, for an average week, excluding promotion activities. The expected demand pattern μ_t is showed in Table 5.1. The data are based on values observed in a practical situation. The values in Table 5.1 are used to determine order-up-to levels corresponding to the Poisson distribution for a target α -service level of 90%. The order-up-to levels are rounded up to integer values. For the MILP model calculations, expected demand data is not rounded. The realisations of demand in the simulations take integer values.

Table 5.1 Expected Poisson demand

	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
Periods t	1	2	3	4	5	6	7
μ_t	3.5	2.3	3.0	2.8	4.5	4.2	2.0

The variable purchasing cost is $c = 1$ per unit and the holding cost is $h = 0.01$ per unit for items that are in the store overnight. The target α -service level is 90% and the maximum shelf life $M = 3$. We vary the fixed setup cost k (0, 1 and 3), the disposal cost w (0, 0.5 and -0.5) and the fraction l of LIFO demand (0, 0.4 and 0.6). A disposal cost of 0 means that in case of waste, only the purchasing cost is lost. A positive disposal cost incurs a cost to dispose of the waste and a negative disposal cost means that there is a salvage value for the wasted items. Varying the parameters k , w and l gives 27 scenarios, for which parameter values for the YS order policy is generated by the MILP model. For each scenario, the four policies introduced in Section 5.4 are simulated during 10,000 successive weeks.

5.6.2 Results of 27 scenarios

Table 5.2 shows the results of the four simulated policies for each scenario. The number of orders the MILP model prescribes are in the column '# Orders'. This number applies to all policies. The average total costs of the simulated policies are depicted relative to the costs of the YS_{2I} policy simulation (100). For every policy, the number of days per week the target α -service level is not reached is given in the columns '# < α '. For every instance, the average fill rate per week, being the fraction of demand fulfilled from stock, is measured.

For the $YS_{\Sigma I}$ simulation, the average fill rate is at least 96%. For YS_{Fresh} , $YS_{\gamma \Sigma I}$ and \hat{YQ}_{rt} , the average fill rates per week are at least 90%, 98% and 96% respectively.

Table 5.2 Overview of four simulated policies to calculate the order quantity

Scen	k	LIFO	w	$YS_{\Sigma I}$		YS_{Fresh}		$YS_{\gamma \Sigma I}$		\hat{YQ}_{rt}	
				# Orders	# < α	TotCost	# < α	TotCost	# < α	TotCost	# < α
1	0	FIFO	0	7		106.3		103.4		109.7	
2			0.5	7		108.3		104.2		112.7	
3			-0.5	7		104.2		102.5		106.4	
4		0.4	0	7	1	109.0		104.5		109.1	
5			0.5	7	1	111.8		105.6		111.8	
6			-0.5	7	1	105.9		103.2		106.1	
7		0.6	0	7	1	110.4		106.9		108.2	
8			0.5	7	1	113.6		108.9		110.7	
9			-0.5	7	1	106.8		104.6		105.4	
10	1	FIFO	0	4		106.0	1	104.9		111.4	
11			0.5	5	1	105.6	1	105.9		117.4	
12			-0.5	4		102.5	1	103.3		107.0	
13		0.4	0	7	1	107.4		104.4		109.1	
14			0.5	7	1	109.7		105.3		111.2	
15			-0.5	4	1	102.4	1	104.5		102.2	1
16		0.6	0	6	3	110.3		107.4		109.3	1
17			0.5	7	1	112.3		108.2		110.9	
18			-0.5	4	3	104.3	1	105.4	2	106.2	
19	3	FIFO	0	4		104.2	1	103.9		109.0	
20			0.5	4		106.8	1	105.1		112.2	
21			-0.5	3		96.4	2	103.5		107.9	
22		0.4	0	3	2	99.5	2	106.4		106.9	
23			0.5	4	1	106.8	1	106.9		102.9	1
24			-0.5	3	2	96.5	2	104.2		104.6	
25		0.6	0	3	2	98.5	2	109.3		105.1	
26			0.5	4	3	108.5	1	108.7	2	110.1	
27			-0.5	3	2	96.1	2	106.2		103.9	
Average						105.6		105.4		108.4	

As can be observed from Table 5.2, the $YS_{\Sigma I}$ policy gives mostly the lowest cost solution, but does not always meet the α -service level requirement. However, the lowest observed reached α -service level in case of FIFO demand is 88%, which is only observed for one instance. In case of 0.4 LIFO demand, 83% is the lowest reached α -service level and 79% at 0.6 LIFO demand. The performance of the YS_{Fresh} policy is poor. In case $k = 3$, in 5 scenarios the YS_{Fresh} policy has the lowest costs, but in two days of the week the α -service level requirements are not met, with lowest values up to 31% in scenario 25 and 27. The

YS_{Fresh} policy often gives an infeasible solution, and when the required α -service levels are met, there is a cheaper feasible solution available from the other investigated policies. The $YS_{\gamma SI}$ policy is used with a fraction of $\gamma = 0.8$ for FIFO demand, such that all FIFO scenarios have a feasible solution. For most scenarios this fraction could even be higher, but we did not search for the best value of γ per scenario. For $l = 0.4$ LIFO demand we used a fraction of $\gamma = 0.78$ and for $l = 0.6$ LIFO demand we used a fraction of $\gamma = 0.72$. In the latter case, the solutions of scenarios 18 and 26 are not completely feasible. To reach that, γ should be lowered to 0.24. The $YS_{\gamma SI}$ policy fits to the situation in retail practice where the age-distribution of the items on the shelf is unknown. Generally it can be tuned towards feasible solutions and the average total costs are higher than in the YS_{SI} policy but lower than the other investigated policies. Finally, the $Y\hat{Q}_{rt}$ policy gives feasible solutions with three exceptions where an α -service level of 87%, 87% and 89% is reached. This procedure is generally more costly than the other policies.

5.6.3 Scenario in detail

To get insight in the inventory dynamics in the situation where not every day can be ordered and the LIFO part of demand is $l = 0.4$, we investigate scenario 22 in more detail, since the YS_{SI} policy does not give a feasible policy for this scenario. The fixed setup cost is $k = 3$ and the cost of waste is $w = 0$. Table 5.3 shows the order-up-to level, the expected order quantity and the expected waste for every day of the week, as calculated with the MILP model. The average values of the simulation and the reached α -service levels are depicted in the other rows. The simulated order quantities are a bit higher as well as the expected waste. The required α -service level is not met on Tuesday and Thursday.

Simulation of the YS_{Fresh} policy gives order quantities that deviate considerably from that of the YS_{SI} policy. The system behaves more nervous if the determination of the order quantity only considers the freshest items. Furthermore, on Thursday the α -service level reaches only 34%. The average total costs are only 0.5% lower than those of the YS_{SI} policy. Simulation of the $YS_{\gamma SI}$ policy shows, naturally, higher order quantities than the YS_{SI} policy, but the system is less nervous than in the YS_{Fresh} policy. The average total costs are 6.4% higher than the average total costs of the YS_{SI} policy. The α -service level requirement is met every day. The $Y\hat{Q}_{rt}$ policy resembles the $YS_{\gamma SI}$ approach in this case.

When the fixed setup cost $k = 0$, but not every day can be ordered, one can also put a limit on the number of orders per week in the MILP model. Setting $\Sigma Y_t = 3$ gives the same solution as presented in Table 5.3 with $k = 3$. Setting $\Sigma Y_t = 4$ leads to orders on Tuesday, Thursday, Friday and Sunday. Another possibility is to fix the days to order and determine the order-up-to levels with the MILP model.

Table 5.3 Scenario $k = 3$, $w = 0$ and LIFO = 0.4 in detail for all policies and results for a double expected demand

	t	Mon	Tues	Wed	Thurs	Fri	Sat	Sun	Total
	μ_t	3.5	2.3	3.0	2.8	4.5	4.2	2.0	22.3
MILP	policy S	0	13.40	0	19.20	0	0	15.50	
	EQ	0	7.90	0	12.50	0	0	9.00	29.40
	Ewaste	0	0	1.40	0	1.20	0	4.50	7.10
$YS_{\Sigma I}$	avg Q	0	7.87	0	13.45	0	0	8.88	30.20
	avg waste	0	0	2.18	0	1.60	0	4.80	8.58
	reached α	0.96	0.85	1.00	0.88	1.00	0.98	0.92	
YS_{Fresh}	avg Q	0	2.91	0	13.29	0	0	14.67	30.88
	avg waste	0	0	6.77	0	0.38	0	3.51	10.66
	reached α	1.00	0.98	1.00	0.34	1.00	0.95	0.74	
$YS_{\gamma \Sigma I}$ $\gamma = 0.78$	avg Q	0	8.79	0	14.29	0	0	9.74	32.82
	avg waste	0	0	2.66	0	2.23	0	5.88	10.77
	Alpha	0.99	0.90	1.00	0.95	1.00	0.99	0.93	
\hat{YQ}_{rt}	avg Q	0	9	0	15	0	0	9	33.00
	avg waste	0	0	1.92	0	2.26	0	6.68	10.86
	reached α	1.00	0.91	1.00	0.97	1.00	1.00	0.98	
	$2 \cdot \mu_t$	7.0	4.6	6.0	5.6	9.0	8.4	4.0	44.6
MILP	policy S	0	22.80	0	20.00	27.00	0	22.40	
	EQ	0	12.80	0	9.60	12.60	0	12.80	47.8
	Ewaste	1.4	0	1.8	0	0	0	0	3.2
$YS_{\Sigma I}$	avg Q	0	13.04	0	8.60	9.54	0	10.89	42.07
	avg waste	4.68	0	5.57	0	6.26	0	3.35	19.86
	reached α	0.99	0.95	1.00	0.99	1.00	1.00	0.99	

Additionally, we investigate the case when the expected demand is twice as high and the other parameters are kept constant. Table 5.3 shows the MILP solution and the simulated $YS_{\Sigma I}$ results. As can be observed, the α -service level requirement is met on all days of the week.

In the MILP model the expected amount of freshest items is overestimated and the expected waste is underestimated, as explained in Chapter 2 (Pauls-Worm et al., 2014). The small expected demand numbers in the Poisson distribution in a retail situation enlarge this effect compared to a situation with a high expected demand. In case of a higher expected demand, the performance of the MILP generated parameter values and the $YS_{\Sigma I}$ order policy to calculate the order quantity improves.

5.7 Conclusion

We studied retail order policies for a practical situation where stores may be able to order every day, or on fixed days of the week. We investigated the inventory system of a perishable product with a remaining shelf life on delivery at the store of three days and a

weekly (seasonal) demand pattern which is stationary over the weeks. The stores have a target α -service level. Customer demand is either FIFO or combined LIFO – FIFO, with a LIFO fraction of 0.4 or 0.6. Lead time is one day. An SP model of the situation in the retailer practice has been presented. Moreover, we derived an MILP approximation model according to the static-dynamic uncertainty strategy, which generates parameter values for a YS order policy for a perishable product under an α -service level constraint with a lead time of one day. For the practical application of the YS policy we developed four policies that determine the order quantity based on the YS parameter values found by the MILP model. In the $YS_{\Sigma I}$ policy the order quantity is determined considering the total available inventory. This gives lowest cost solutions, which are not always feasible. From the three other policies, $YS_{\gamma \Sigma I}$ where a fraction of the total available inventory is considered when determining the order quantity, performs best. One can always find a fraction that gives feasible solutions. This policy fits to the situation in retail practice where the age-distribution of the items on the shelf is unknown. The calculations take less than a second and can be executed with standard software, which is important for implementation in practice. We showed that the developed approach can find a suitable reorder plan according to a static-dynamic uncertainty strategy where the order days are fixed.

Chapter 6

MILP approximation generating a YQ policy

Based on:

Pauls-Worm, K.G.J., Hendrix, E.M.T., Alcoba, A.G., Haijema, R., 2015. *Order quantities for perishable inventory control with non-stationary demand and a fill rate constraint*. International Journal of Production Economics xxx, pp. xxx-xxx.

Abstract

We study the practical decision problem of fresh food production with a long production lead time to decide every period (e.g. week) how many items to produce. When a batch is ready for use, its items have a fixed shelf life, after which the items become waste in the sense that they cannot be sold anymore. The demand for (fresh) food products is uncertain and highly fluctuating, mainly caused by price promotions of retail organisations. We focus on cases where a so-called cycle fill rate service level requirement applies. We investigate the generation of a production plan that fixes the timing and quantity of the production for a finite time horizon. To minimise waste, one issues the oldest items first, i.e. a FIFO issuing policy. In case of out-of-stock, sales are lost.

We model this decision problem as a Stochastic Programming (SP) model. The objective of our study is to find order quantities for the SP model, that approximately meet cycle fill rate service level requirements while keeping outdating low. To find approximate solutions for the SP model, an MILP model is developed. The MILP model is a deterministic approximation that generates feasible replenishment quantities in less than a second. With a scenario-based MINLP approach, optimal solutions are generated for a large sample of demand paths as a benchmark for the MILP solutions. We show that the MILP model is suitable for practical use if the setup cost is such that the replenishment cycles in the production plan are close to or of the same length as the maximum shelf life. In those cases, the expected total costs are close to the costs of the optimal solution and the average fill rate is close to the required one.

6.1 Introduction

A producer of a fresh food product with a long lead time has to decide every period (e.g. week) how many items to produce. Multiple batches of the same product of different ages are in production, so there are multiple outstanding orders. When a batch is ready for use, the items get a 'best before' or a 'use by' date, resulting in a fixed shelf life for the product. Generally, producers have a contract with retail organisations about the minimum remaining shelf life of the items delivered. The time between 'ready for use' and the minimum remaining shelf life is the maximum internal shelf life the producer can use to organise production efficiently. After the maximum internal shelf life of several periods, the product becomes waste in the sense that it cannot be sold anymore with the aimed remaining shelf life. It still may be used for other purposes, so it may have a salvage value. These production characteristics can be found in e.g. the maturation of cheese, meat from breeding to slaughtering and crops from seed to harvesting. In the remainder of this chapter we will use the shorter term 'maximum shelf life' to denote the maximum internal shelf life.

The demand for (fresh) food products is uncertain, so the production quantity is determined based on forecasts. A complicating factor is that demand is highly fluctuating, mainly due to price promotions of the retail organisations, i.e. demand is non-stationary. Competition in retail is very strong, so retail organisations are reluctant to share information about their promotional activities. Because of the long lead time, planned promotions sometimes become only known to the food producer after production of the items. However, retail organisations continuously work on improving their demand forecasts. Highly fluctuating demand is not necessarily highly uncertain. It is likely that the food producer has a contract with the retail organisation about the delivery of the product, with respect to remaining shelf life and fill rate. The fill rate indicates that a predefined percentage of the demand per replenishment cycle has to be fulfilled from stock, a so-called cycle fill rate. According to the food producer, demand that cannot be fulfilled from stock is lost. The food producer has control over the issuing of the items. In order to minimise waste due to outdating, the oldest items are issued first, so there is a first in – first out (FIFO) issuing policy. The food producer wants to make a production plan minimising the expected total costs, indicating when to produce and how much, for a fixed time horizon of T weeks.

Having a long lead time, the inventory levels at the time of delivery are unknown at the time the replenishment quantity has to be determined. So the replenishment quantity should be decided on beforehand, independently of the inventory level I . Therefore, we consider a so-called static uncertainty policy denoted by YQ , where Y denotes in which periods to deliver, and Q the corresponding replenishment quantities to deliver at the beginning of period t . The research question is whether it is possible to generate a production plan for T periods using existing solvers and for which instances the solution might be close to the optimal solution.

The rest of the chapter is organised as follows. In Section 6.2, we explain how the chapter contributes to literature by discussing the main characteristics of the problem. In Section 6.3, the problem is formulated as a Stochastic Programming (SP) model. Section 6.4 describes the steps towards a production plan that we use to formulate a deterministic MILP model that generates feasible production plans. The model determines the timing of deliveries, the replenishment cycle length and the replenishment or delivery quantity, in order to make a production plan for lead time L periods before t . This model is presented in Section 6.5. Section 6.6 investigates the applicability of the MILP model compared to 'optimal' solutions generated by a scenario-based MINLP approach. Section 6.7 concludes and provides topics for future research.

6.2 Literature

The lost sales inventory problem studied in this chapter is to fix a production plan (and consequently replenishment plan) for a single perishable product with a long lead time, non-stationary demand and a fill rate constraint. The chapter builds upon previous studies by Bookbinder and Tan (1988), Tarim and Kingsman (2004) and Pauls-Worm et al. (2014) (Chapter 2), which we will describe in more detail below. This chapter contributes to these studies by moving from a zero lead time problem, where the current inventory levels can be taken into account, to a problem with long lead time. Moreover, motivated by a practical case, the problem includes a fill rate or β -service level constraint instead of an α -service level constraint. Consequently, also a different order policy is studied. In Chapter 2 (Pauls-Worm et al., 2014), the derived order policy is of the type $Y_t S_t$, that is, in period t , if $Y_t = 1$ the manager orders up to a level S_t . Due to the long lead time, in the present chapter we consider a $Y_t Q_t$ policy; if $Y_t = 1$ in period t , a fixed quantity Q_t will be delivered and consequently produced the lead time L periods before t .

Bookbinder and Tan (1988) formulated an SP model for a single-item inventory problem for a finite horizon, with a non-stationary demand, under an α -service level constraint. They distinguish two decision rules, a static uncertainty and a dynamic uncertainty strategy. In the static uncertainty strategy the timing and order sizes $Y_t Q_t$ are determined at the beginning of the time horizon, before demand is known. A dynamic uncertainty strategy bases decisions on new available information. Bookbinder and Tan (1988) combine the two strategies into a third static-dynamic uncertainty $Y_t S_t$ policy. Tarim and Kingsman (2004) formulated for a non-perishable product an MILP model for the static-dynamic uncertainty strategy resulting in an optimal $Y_t S_t$ policy. Rossi (2013b) assessed the quality of a Constraint Programming solution procedure on a static-dynamic $Y_t S_t$ policy for perishable items. Pauls-Worm et al. (2014) extended in Chapter 2 the model of Tarim and Kingsman (2004) to perishable items, resulting in an approximate solution for the SP model for lead time zero and an α -service level.

For the static uncertainty $Y_t Q_t$ policy and non-perishable items with a non-stationary demand, Tempelmeier and Herpers (2011) formulated a Stochastic Single Item

Uncapacitated Lot-Sizing Problem with a fill rate constraint. They assume that excess demand is backlogged and they found an optimal $Y_t Q_t$ solution with a modification of Dijkstra's shortest-path algorithm, as well as several heuristic approaches to solve the model.

Key characteristics of the lost sales inventory problem under study are perishability, non-stationary demand, fill rate constraints, and long lead time. In the remainder of this section, we discuss how these aspects are addressed in literature.

Perishability and non-stationary demand

Regarding perishability, we focus on products with a fixed shelf life. A recent review about perishable inventory, including products with a fixed lifetime, is due to Karaesmen et al. (2011). From their review, it becomes clear that inventory problems with a fixed shelf life, non-stationary demand, and a (long) lead time are relevant but challenging and understudied. Most studies on ordering products with a fixed shelf life, focus on stationary demand. Bijvank and Vis (2011) reviewed lost-sales inventory theory. For non-perishables, they conclude more research should focus on non-stationary demand. Tunc et al. (2011) discuss the use of stationary inventory policies when demand is non-stationary for non-perishable products. They conclude that in case of high demand variability, using a stationary policy can be very expensive. In case of high uncertainty, high setup cost and low penalty cost, using a stationary policy might be efficient. The above papers motivate the interest in policies for non-stationary demand.

A few articles have been published that deal with perishability and non-stationary demand. An exact method to solve the non-stationary problem is Stochastic Dynamic Programming (SDP). For a lead time of one period and a fixed shelf life of up to 7 periods, Haijema et al. (2007) and Haijema et al. (2009) solve the non-stationary ordering problem by SDP and discuss the near optimality of a periodic review order-up-to S_t policy. Their problem, however, lacks a service level constraint. Instead they apply a cost structure that includes a penalty for lost sales. In these studies, fixed setup cost can be included and order periods may be prefixed instead of being part of the optimisation. Some heuristic approaches are published, like the one of Broekmeulen and Van Donselaar (2009). They develop a heuristic for a single store to determine a replenishment policy based on the estimated withdrawal and aging of the items in stock. Similar to Haijema et al. (2007), the weekly demand is stationary with a non-stationary demand pattern during the week. They do not consider fixed setup cost and use a lost sales cost to influence the fill rate.

Service levels

The inclusion of α -service level or fill rate constraints in an optimisation model such as SDP is complicated and subtle. Chen and Krass (2001) define the difference between mean service level constraints and minimal service level constraints. A mean service level constraint measures the service level over the time horizon, while a minimal service level constraint measures the service level in every period. We use a minimal service level criterion, a minimal fill rate per replenishment cycle. Food production companies often have service contracts with their retail customers requiring a certain fill rate service level.

According to Chen and Krass (2001) a minimal service level criterion is preferred when the service level constraint is due to a contractual obligation or a company policy. Minner and Transchel (2010) determine a replenishment policy for perishable products in retail assuming a weekly demand pattern, negligible fixed cost, and positive but relatively short lead times. They apply an SDP model with marginal α -service level and fill rate constraints. What they call marginal service levels are in terms of Chen and Krass (2001) minimal service levels as they should hold per (sub)period. Hendrix et al. (2012) also apply a minimal α -service level in an SDP approach for perishable products. Note that such an approach generates a $Q(I)$ policy (a dynamic uncertainty rule), which is not suitable in cases of long lead time. In Chapter 3, Pauls-Worm and Hendrix (2015) show when considering service level constraints, SDP generates order policies that are not necessarily optimal. SDP meets a service level requirement that is conditional to each possible starting inventory level, no matter how small the chance of occurrence. This results in an overachievement of the service level constraint.

Lost sales and long lead time

Base stock policies are commonly studied policies for having appropriate structural properties and being close to optimal in many settings, especially in cases of backlogging and short lead times. Morton (1969) showed for non-perishables with a stationary demand that an order-up-to policy is not optimal for a lost-sales inventory model in case of a positive lead time. Van Donselaar et al. (1996) show for non-perishables, that compared to a replenishment policy with static order-up-to levels, it may be more efficient in a lost sales system to use dynamic order-up-to levels that dynamically meet fill rate constraints. Under dynamic order-up-to levels, the pattern of successive order sizes is smoother. The need for smoother order patterns is even stronger in case of long lead times, as shown in Goldberg et al. (2014). They show that as the lead time grows for non-perishables with a stationary demand, the constant-order policy is asymptotically optimal. The intuition of their approach is to select the constant-order policy that considers the inventory in the pipeline best. In case of non-stationary demand, the expected pipeline inventory is highly fluctuating per period because of the fluctuating demand. The accuracy of approximated pipeline inventory will be low. In this chapter we deal with non-stationary demand of perishable items, meaning that at the end of the shelf life items will become waste. This results in an even less accurate approximation of the pipeline inventory. In the model for long lead times, presented in the next section, we thus study a policy with fixed replenishment quantities Q_t that are independent of the inventory available at the beginning of each replenishment period t .

6.3 Stochastic Programming Model

To determine a production plan for a perishable product with a long lead time under a fill rate constraint, we consider a single-product – single-echelon SP model, minimizing expected total costs. We focus on an $Y_t Q_t$ policy. Periods are of equal lengths and can be hours, days, weeks or months, whatever is applicable in the practical situation. The product

has a fixed maximum integer (internal) shelf life $M \geq 2$ periods. Due to the findings of Goldberg et al. (2014), we leave lead time out of the model. The model solution is concerned with the delivery time t of a production batch, i.e. production or ordering should be done L periods before delivery. Demand is non-stationary independently distributed with a Normal distribution $d_t \sim N(\mu_b (CV\mu_t)^2)$ in period t . We use a Normal distribution to keep fill rate calculations simple. Demand is never negative, food cannot be returned due to safety regulations. There are fixed and variable production costs, holding cost and cost of waste. Table 6.1 shows the list of symbols. Fill rate is defined as the proportion of demand per replenishment cycle that can be fulfilled directly from stock, being a cycle fill rate. The (maximum) available inventory is determined by the replenishment quantity, which should be such that the predetermined fill rate constraint can be met. Issuing is according to a FIFO policy in which the first delivered items are issued first. The age of the items is indexed by $b = 1, \dots, M$. The inventory level of age b at the end of period t is denoted by I_{bt} . Items delivered at the beginning of period t , have age $b = 1$ at the end of period t . Items of age M at the end of the period are considered waste and cannot be used in the next period. If demand in period t exceeds the available inventory of period $t - 1$ plus the delivered quantity Q_t at the beginning of t , demand is lost and there is a shortage of X_t . Replenishments can aim to cover demand from 1 up to M periods, so replenishment cycles have a varying length j .

Table 6.1 List of symbols

T	length of finite time horizon
t	period index , $t = 1, \dots, T$
M	fixed maximum (internal) shelf life
b	age index, $b = 1, \dots, M$
j	index denoting the length of the replenishment cycle
k	fixed setup cost for every replenishment
c	variable unit production cost
h	unit holding cost, for items that are carried over from one period to the next
w	unit disposal cost ($w > 0$) or salvage value ($w < 0$) for items becoming waste
β	target cycle fill rate
d_t	stochastic demand during period t , non-stationary for $t = 1, \dots, T$
Y_t	binary variable takes the value of 1 if there is a replenishment in period t , and 0 otherwise
Q_t	replenishment quantity for delivery at the beginning of period t
I_{bt}	inventory level of items with age b at the end of period t
I_{Mt}	inventory of age M at the end of period t is considered waste
X_t	number of items short at the end of period t
$E(TC)$	expected total costs over the time horizon

This problem is formulated as a stochastic programming model.

$$\text{Min } E(TC) = \sum_{t=1}^T \left(kY_t + cQ_t + E \left(h \sum_{b=1}^{M-1} I_{bt} + wI_{Mt} \right) \right) \quad (1)$$

subject to

$$Y_t = \begin{cases} 1 & \text{if } Q_t > 0 \\ 0 & \text{otherwise} \end{cases} \quad t = 1, \dots, T \quad (2)$$

$$I_{bt} = \left(I_{b-1,t-1} - \left(d_t - \sum_{i=b}^{M-1} I_{i,t-1} \right)^+ \right)^+ \quad t = 1, \dots, T; \quad b = 2, \dots, M \quad (3)$$

$$I_{1t} = \left(Q_t - \left(d_t - \sum_{b=1}^{M-1} I_{b,t-1} \right)^+ \right)^+ \quad (4)$$

$$X_t = \left(d_t - \sum_{b=1}^{M-1} I_{b,t-1} - Q_t \right)^+ \quad t = 1, \dots, T \quad (5)$$

$$Y_t \cdot \sum_{j=1}^{\min\{M, T-t+1\}} \left(\frac{\sum_{i=t}^{t+j-1} E(X_i)}{\sum_{i=t}^{t+j-1} \mu_i} \cdot Y_{t+j} \cdot \prod_{i=t+1}^{t+j-1} (1 - Y_i) \right) \leq (1 - \beta) \quad t = 1, \dots, T \quad (6)$$

$$Y_{T+1} = 1 \quad (7)$$

$$I_{b0} = 0 \quad b = 1, \dots, M \quad (8)$$

$$I_{bt} \geq 0 \quad t = 1, \dots, T; \quad b = 1, \dots, M \quad (9)$$

$$Y_t \in \{0, 1\}, \quad Q_t, X_t \geq 0 \quad t = 1, \dots, T \quad (10)$$

The model minimises the expected total costs (Eq. 1), consisting of fixed setup cost for every replenishment, production cost and the expected cost of holding inventory and of waste. Eqs. (3), (4) and (5) model the inventory levels of all ages and shortage in period t , under a FIFO issuing policy. To meet the fill rate requirement, the fraction of expected shortage over expected demand of a replenishment cycle should be less than or equal to $(1 - \beta)$. Index j denotes the length of the replenishment cycle, which has a length of $j = 1$ to M periods. The first period after the replenishment cycle, i.e. period $t + j$, should have a delivery, and during the replenishment cycle no other delivery takes place. Eq. (6) models the fill rate requirement related to the length of the replenishment cycle, where $Y_t = Y_{t+j} = 1$, and $Y_i = 0$ with $t < i < t + j$. Eqs (7) to (10) are definition constraints. The model assumes that at the end of the time horizon T , also the last replenishment cycle ends, so in $T + 1$ a new replenishment arrives (Eq. 7). In the evaluations, the starting inventory level is zero (Eq. 8).

6.4 Towards a production plan

Eq. (6) describes the fill rate requirement $\frac{\sum_{i=t}^{t+j-1} E(X_i)}{\sum_{i=t}^{t+j-1} \mu_i} \leq (1 - \beta)$ for all replenishment cycle

lengths j . Consider the replenishment quantity Q_t for which this constraint holds when the inventory at delivery period t is zero, i.e. $\sum_{b=1}^{M-1} I_{b,t-1} = 0$. The aggregated demand

$d = \sum_{i=t}^{t+j-1} d_i$ during the replenishment cycle of length j has a normal distribution with mean $\mu = \sum_{i=t}^{t+j-1} \mu_i$ and variance $\sigma^2 = \sum_{i=t}^{t+j-1} \sigma_i^2 = \sum_{i=t}^{t+j-1} (\text{CV} \cdot \mu_i)^2$. The expected shortage

$E(X) = \sum_{i=t}^{t+j-1} E(X_i)$ is the so-called loss function, expressing the expected shortage as a function of the delivered quantity Q_t . Let ϕ be the density function (pdf) and Φ the cumulative distribution function (cdf) of d . Then the loss function is

$$L(Q_t) = E(X) = E(d - Q_t)^+ = \int_Q^\infty (x - Q_t) \phi(x) dx. \quad (11)$$

It is known that the loss function is convex in quantity Q_t and for the normal distribution can be expressed by (Chopra and Meindl, 2010)

$$L(Q_t) = \mu - Q_t + \sigma \cdot \left(\phi\left(\frac{Q_t - \mu}{\sigma}\right) + \Phi\left(\frac{Q_t - \mu}{\sigma}\right) \cdot \frac{Q_t - \mu}{\sigma} \right). \quad (12)$$

The minimum replenishment quantity Q_t fulfilling the fill rate requirement can be found by solving $L(Q_t) = (1 - \beta)\mu$, i.e.

$$Q_t - \sigma \cdot \left(\phi\left(\frac{Q_t - \mu}{\sigma}\right) + \Phi\left(\frac{Q_t - \mu}{\sigma}\right) \cdot \frac{Q_t - \mu}{\sigma} \right) = \beta\mu. \quad (13)$$

We solved Eq. (13) using a standard solver "fzero" of MATLAB. As shown in Alcoba et al. (2015), for a replenishment cycle of $j = 1$ period and the defined distribution, it is sufficient to solve Eq. (13) for all periods t . A replenishment cycle can have a length of $j = 1, 2, \dots, M$ periods. An $M \times T$ table called *Level/Q* can be generated with all possible replenishment quantity levels *Level/Q_{jt}* for period t and replenishment cycle length j when the inventory is

zero, i.e. $\sum_{b=1}^{M-1} I_{b,t-1} = 0$. Table 6.2 shows an example of the expected demand per period,

and the corresponding *Level/Q_{jt}* values.

Table 6.2 Expected demand μ_t and corresponding $Level/Q$ values for $CV=\sigma/\mu=0.25$ and fill rate $\beta=95\%$.

t	1	2	3	4	5	6	7	8	9	10	11	12
μ_t	800	950	200	900	800	150	650	800	900	300	150	600
$Level/Q_{1t}$	899	1068	225	1011	899	169	731	899	1011	337	169	674
$Level/Q_{2t}$	1832	1243	1187	1779	1030	863	1518	1779	1280	475	807	0
$Level/Q_{3t}$	2011	2114	1958	1913	1652	1652	2390	2051	1414	1085	0	0

In the next section, we will show how the $Level/Q_{jt}$ values can be used to find a $Y_t Q_t$ policy with an MILP model.

6.5 MILP model

In Section 6.5.1, we formulate an MILP model to generate approximate solutions for the SP model. In Section 6.5.2, a numerical illustration of the MILP model is presented.

6.5.1 MILP model formulation

Besides the policy variables Y_t and Q_t , the other variables of the MILP model are denoted by their expected value variant, i.e. EI and EX , in contrast to the SP model. The objective function is given by

$$\text{Min } E(TC) = \sum_{t=1}^T \left\{ kY_t + cQ_t + h \sum_{b=1}^{M-1} EI_{bt} + wEI_{Mt} \right\} \quad (14)$$

and minimises the setup cost, the inventory holding cost over the on-hand inventory at the end of the period, the variable production cost and the cost of waste.

$$\sum_{j=1}^M Z_{jt} = Y_t \quad t = 1, \dots, T - M + 1 \quad (15)$$

$$\sum_{j=1}^{T-t+1} Z_{jt} = Y_t \quad t = T - M + 2, \dots, T \quad (16)$$

$$Z_{1t} \leq Y_{t+1} \quad t = 1, \dots, T - 1 \quad (17)$$

$$j \cdot Z_{jt} \leq \sum_{i=1}^{j-1} (1 - Y_{t+i}) + Y_{t+j} \quad t = 1, \dots, T - j; \quad j = 2, \dots, M \quad (18)$$

$$Y_{T+1} = 1 \quad (19)$$

Eqs. (15) – (19) are logical constraints to describe the order timing, thus linearising Eq. (6) of the SP model. Variable $Y_t = 1$ if there is a delivery in period t , and variable $Z_{jt} = 1$ denotes the replenishment cycle length j in period t aimed at fulfilling demand for j periods: for period t and the next $j - 1$ periods. In case of a delivery, eqs. (15) and (16) require a delivery for 1 or 2 or... up to M periods. Eq. (16) is valid at the end of the horizon. If there is no delivery in period t , then $Z_{jt} = 0$ for all j . On the other hand, a delivery in period t

implies the choice of exactly one replenishment cycle length j , i.e. Z_{jt} has to be 1 for one value of j . For replenishment cycle length j (covering t to $t + j - 1$), a new replenishment takes place in period $t + j$ ($Y_t = 1$ and $Y_{t+j} = 1$) and in between there is no order (equations (17) – (19)).

$$Y_1 = 1 \quad (20)$$

$$\sum_{j=1}^M Y_{t+j-1} \geq 1 \quad t = 1, \dots, T - M \quad (21)$$

Eqs. (20) and (21) are constraints to ensure that at least in the first period and every M periods an order is delivered.

$$Q_t \geq \sum_{j=1}^M (Z_{jt} \cdot \text{Level}Q_{jt}) \quad t = 1, \dots, T - 1 \quad (22)$$

Eq. (22) selects replenishment quantity Q_t for j periods from the table $\text{Level}Q_{jt}$ (Section 6.4). The model ignores the pipeline inventory. When the replenishment cycle is of length M , the pipeline inventory is zero. In other cases the replenishment is greater than strictly necessary to meet the cycle fill rate requirement, causing an approximate solution of the formulated SP model. The size of the pipeline inventory when replenishment cycles are of length $< M$ depends on the distribution of demand in successive periods and is therefore hard to approximate.

$$EI_{M-1,t-1} - Ed_t = EI_{Mt} - EA_{M-1,t} \quad t = 1, \dots, T \quad (23)$$

$$EI_{b,t-1} - EA_{b+1,t} = EI_{b+1,t} - EA_{bt} \quad t = 1, \dots, T; \quad b = 1, \dots, M - 2 \quad (24)$$

$$Q_t - EA_{1t} = EI_{1t} - EX_t \quad t = 1, \dots, T \quad (25)$$

Eq. (23), (24) and (25) keep track of the age-distribution of the items in stock, under a FIFO-issuing policy. Let auxiliary variable EA_{bt} denote the shortage of inventory of age b with $b = 1, \dots, M - 1$ in period t to fulfil the demand of period t . If EA_{bt} has a positive value, then fresher inventory is used to fulfil demand. Eq. (23) imposes the oldest inventory to be used first to fulfil demand. What is left over has the maximum shelf life and will become waste, or shortage of the oldest inventory occurs. In the latter case, Eq. (24) is appropriate. The shortage has to be fulfilled by items of intermediate ages, until demand is fulfilled by the freshest items that have been delivered in the current period, according to Eq. (25). The FIFO constraints linearise constraints (3) – (5) of the SP model, causing an over- and underestimation of the different inventory levels as shown in Chapter 2 (Pauls-Worm et al., 2014). Chapter 2 also shows that the FIFO constraints are necessary to meet FIFO issuing in the context of an MILP model.

$$\mathcal{M}_t \cdot BA_{bt} \geq EA_{bt} \quad t = 1, \dots, T; \quad b = 1, \dots, M - 1 \quad (26)$$

$$\mathcal{M}_t \cdot (1 - BA_{bt}) \geq EI_{b+1,t} \quad t = 1, \dots, T; \quad b = 1, \dots, M - 1 \quad (27)$$

$$\mathcal{M}_t \cdot BX_t \geq EX_t \quad t = 1, \dots, T \quad (28)$$

$$\mathcal{M}_t \cdot (1 - BX_t) \geq EI_{1t} \quad t = 1, \dots, T \quad (29)$$

At most one variable at the right-hand-sides of equations (23), (24) and (25) can have a positive value. The other variable needs to have a value of 0. Equations (26) – (29) take

care of that, using the binary variables BA_{bt} and BX_t . M_t is a sufficiently large number, for instance $M_t = Level/Q_{Mt}$.

$$EI_{b0} = 0 \quad b = 1, \dots, M \quad (30)$$

The starting inventory is zero (Eq.(30)).

$$EI_{bt}, Q_t \geq 0 \quad t = 1, \dots, T; \quad b = 1, \dots, M \quad (31)$$

$$EA_{bt} \geq 0 \quad t = 1, \dots, T; \quad b = 1, \dots, M - 1 \quad (32)$$

$$EX_t \geq 0 \quad t = 1, \dots, T \quad (33)$$

$$Y_t, Z_{jt} \in \{0, 1\} \quad t = 1, \dots, T; \quad j = 1, \dots, M \quad (34)$$

$$BA_{bt}, BX_t \in \{0, 1\} \quad t = 1, \dots, T; \quad b = 1, \dots, M - 1 \quad (35)$$

Eqs. (31) to (35) are definition constraints.

6.5.2 Numerical illustration of the MILP model

We consider a base case for the MILP model with a fixed setup cost $k = 500$, inventory holding cost $h = 0.5$ over the on-hand inventory at the end of the period, variable production cost $c = 2$, cost of waste $w = 0$, $CV = \sigma/\mu = 0.25$ and required fill rate $\beta = 95\%$. The maximum shelf life is $M = 3$. Cost of waste $w = 0$ implies that for wasted items, there is no extra cost of disposal, nor a salvage value. However, the production cost and holding cost during $M - 1$ periods are still imposed on these items. Expected demand (repeated in the first row of Table 6.3) and $Level/Q_{jt}$ is given in Table 6.2. The solution of this model is given in Table 6.3. The shaded row shows the replenishment quantities with $Y_t = 1$ if $Q_t > 0$. The production plan prescribes deliveries in periods 1, 4, 7, 9 and 12, resulting in 5 replenishment cycles of respectively length 3, 3, 2, 3 and 1. The replenishment quantities are equal to the corresponding values of $Level/Q_{jt}$ in Table 6.2. The expected inventory levels of all ages of period 9 in the table showcase well that the available older inventory is used before the fresh items. The MILP solution is evaluated in a simulation based on 10,000 samples to measure the expected total costs and the fill rate. The expected total costs of

Table 6.3 Numerical illustration of the base case: $E(TC)_{MILP} = 19846$; $E(TC)_{Sim} = 20013$; Avg $\beta_{sim} = 95.44\%$

t	1	2	3	4	5	6	7	8	9	10	11	12
μ_t	800	950	200	900	800	150	650	800	900	300	150	600
Y_t	1	0	0	1	0	0	1	0	1	0	0	1
Q_t	2011	0	0	1913	0	0	1518	0	1414	0	0	674
EI_{1t}	1211	0	0	1013	0	0	868	0	582	0	0	74
EI_{2t}	0	261	0	0	213	0	0	68	0	282	0	0
EI_{3t}	0	0	61	0	0	63	0	0	0	0	132	0
EX_t	0	0	0	0	0	0	0	0	0	0	0	0
β_{sim}	0	0	95.07	0	0	95.01	0	95.06	0	0	97.02	95.04

the simulation are 8.4% higher than the costs of the MILP solution. The final row of Table 6.3 contains the simulated fill rate β_{sim} . In period 11 the fill rate is higher than required, because the replenishment cycle starting in period 9 has on-hand inventory at the start of the period. The levels of $Level/Q$ are based on no on-hand inventory at the beginning of the period. The fill rate of other replenishment cycles is close to the required value. These observations fit with the design of the MILP model and the definition of $Level/Q_{jt}$.

6.6 Results of the MILP model

To investigate the sensitivity and applicability of the MILP model for different parameter values, a design of experiments is set up, reported in Section 6.6.1. Section 6.6.2 describes the benchmark with the scenario-based MINLP approach we use. In Section 6.6.3 a comparison of the MILP solutions is made with solutions of a MINLP scenario-based approach.

6.6.1 Design of Experiments

The experimental design is similar to the one in Chapter 2 (Pauls-Worm et al., 2014). For the setup cost k we consider values 0, 500, and 1000. The time horizon is $T = 12$ periods. The inventory holding cost is $h = 0.5$ over the on-hand inventory at the end of the period, variable production cost is $c = 2$, and the maximum shelf life is $M = 2, 3$ or 4. This setting is based on values used in practice. We learned from Chapter 2 that $M = 3$ is the most interesting case to study. Table 6.4 shows the design of experiments. The design varies the

Table 6.4 Design of Experiments

Experiment	Demand	k	w	Fill rate β (%)	CV	M
Base	Erratic	500	0	95	0.25	3
1 – 9	Erratic	0	-0.5, 0, 0.5	90, 95, 98	0.10	3
10 – 18	Erratic	0	-0.5, 0, 0.5	90, 95, 98	0.25	3
19 – 27	Erratic	0	-0.5, 0, 0.5	90, 95, 98	0.33	3
28 – 36	Erratic	500	-0.5, 0, 0.5	90, 95, 98	0.10	3
*37 – 45	Erratic	500	-0.5, 0, 0.5	90, 95, 98	0.25	3
46 – 54	Erratic	500	-0.5, 0, 0.5	90, 95, 98	0.33	3
55 – 63	Erratic	1000	-0.5, 0, 0.5	90, 95, 98	0.10	3
64 – 72	Erratic	1000	-0.5, 0, 0.5	90, 95, 98	0.25	3
73 – 81	Erratic	1000	-0.5, 0, 0.5	90, 95, 98	0.33	3
82 – 83	Erratic	500	0	95	0.25	2, 4
84	Err Variant	500	0	95	0.25	3
85	Highly Err	500	0	95	0.25	3
86	Stationary	500	0	95	0.25	3

* including the base case

parameter values systematically for setup cost k , cost of waste w , fill rate β and Coefficient of Variation CV. Negative waste cost represents a salvage value for the wasted items, whereas positive waste cost implies disposal cost. This results in 83 experiments using the same erratic demand pattern due to promotion activities of the customer of the producer, the retail organisation. In variation of the base case, the MILP model is also tested with three other demand patterns, being an erratic variant, a highly erratic demand pattern and a stationary demand (Pauls-Worm et al., 2014) (Chapter 2) as depicted in Figure 6.1. The total expected demand of all patterns is 7200.




Fig. 6.1 Demand patterns

6.6.2 Benchmark with a scenario-based approach

The MILP model is a deterministic approach that generates feasible production plans in less than a second for the performed experiments. The question is whether the approach is suitable to use in practice. To investigate this question, an alternative MINLP approach was implemented similar to the approach discussed by Alcoba et al. (2015). For each (integer) feasible timing $Y = (Y_1, \dots, Y_T)$ the best (continuous) replenishment quantities Q_t are generated by nonlinear programming for scenarios consisting of 50,000 sample demand paths. The objective function and fill rate of a quantity vector Q are approximated simulating the inventory development using the 50,000 demand runs. The quantity vector with the lowest average total costs fulfilling the required fill rate is considered to be the optimal (Y_b, Q_b) policy for the given scenarios. In contrast to the MILP approach, it may be clear that due to the enumeration of delivery timings this approach is not tractable as the number of timing vectors grows exponentially in the number of periods. Although that is not a problem for the executed experiments, it may be a problem when the same model described here is applied to other practical cases.

6.6.3 Experiments

Table 6.5 shows the results of the experiments. For each experiment of Table 6.4, the expected total costs of the MILP model ($E(TC)_{\text{MILP}}$), the simulated expected total costs ($E(TC)_{\text{sim}}$) of the MILP policy, the number of orders of the MILP policy (NrO MILP) and the average fill rate of the MILP policy in the simulation (Avg β_{sim}) are given. These are compared with the expected total costs of the MINLP approach ($E(TC)_{\text{MINLP}}$), the number of orders (NrO MINLP) and the average fill rate (Avg β_{MINLP}) of the MINLP policy. The last column shows $E(TC)_{\text{sim}}$ relative to $E(TC)_{\text{MINLP}}$ (x100%), to show the cost increase if an MILP policy is used instead of the MINLP policy. The average fill rate is the average over the fill rates per replenishment cycle. The base case has a grey shade.

Table 6.5 Overview of the results

Exp	k	CV	β	w	$E(TC)_{\text{MILP}}$	$E(TC)_{\text{sim}}$	NrO MILP	Avg β_{sim}	$E(TC)_{\text{MINLP}}$	NrO MINLP	Avg β_{MINLP}	$\frac{E(TC)_{\text{sim}}}{E(TC)_{\text{MINLP}}}$
1	0	0.10	90	-0.5	13124	13195	12	91.29	13050	12	90.17	101.11
2			90	0	13124	13195	12	91.29	13050	12	90.17	101.11
3			90	0.5	13124	13195	12	91.29	13050	12	90.17	101.11
4			95	-0.5	14142	14553	12	97.27	13980	12	95.10	104.10
5			95	0	14142	14553	12	97.27	13980	12	95.10	104.10
6			95	0.5	14142	14553	12	97.27	13980	12	95.10	104.10
7			98	-0.5	16042.5	16112	8	98.78	14934	12	98.05	107.89
8			98	0	16104.5	16238	10	99.16	14935	12	98.05	108.72
9			98	0.5	16165.5	16319	10	99.16	14936	12	98.05	109.26
10		0.25	90	-0.5	14400	15758	12	95.59	13912	12	90.07	113.27
11			90	0	14400	15807	12	95.59	13923	12	90.07	113.53
12			90	0.5	14400	15855	12	95.59	13934	12	90.07	113.79
13			95	-0.5	17154.5	17326	7	96.56	15572	12	95.09	111.26
14			95	0	17344.5	17497	6	95.95	15613	12	95.09	112.07
15			95	0.5	17474	17792	5	95.44	15654	12	95.09	113.66
16			98	-0.5	19185	19171	6	98.34	17334	11	98.07	110.60
17			98	0	19630	19669	6	98.34	17474	12	98.08	112.56
18			98	0.5	20075	20167	6	98.34	17585	12	98.08	114.68
19		0.33	90	-0.5	15923	16479	8	93.85	14604	12	90.07	112.84
20			90	0	15923	16670	8	93.85	14648	12	90.07	113.80
21			90	0.5	15923	16861	8	93.85	14691	12	90.07	114.77
22			95	-0.5	18356.5	18375	5	95.53	16546	12	95.10	111.05
23			95	0	18666	18832	5	95.53	16657	12	95.10	113.06
24			95	0.5	18975.5	19289	5	95.53	16768	12	95.10	115.03
25			98	-0.5	20734	20724	6	98.37	18595	11	98.09	111.45
26			98	0	21430.5	21472	6	98.37	18835	11	98.09	114.00
27			98	0.5	22054	22141	5	98.24	19069	12	98.08	116.11
28	500	0.10	90	-0.5	16875.5	16933	5	90.11	16925	5	90.07	100.05

Exp	k	CV	β	w	$E(TC)_{MILP}$	$E(TC)_{sim}$	NrO		$E(TC)_{MINLP}$	NrO	Avg	$\frac{E(TC)_{sim}}{E(TC)_{MINLP}}$
							MILP	Avg β_{sim}		MINLP	β_{MINLP}	
29			90	0	16875.5	16939	5	90.11	16930	5	90.07	100.05
30			90	0.5	16875.5	16945	5	90.11	16936	5	90.07	100.05
31			95	-0.5	18088	18109	5	95.17	18062	6	95.03	100.26
32			95	0	18088	18145	5	95.17	18088	6	95.03	100.32
33			95	0.5	18088	18181	5	95.17	18114	6	95.03	100.37
34			98	-0.5	19188	19180	5	98.19	19043	6	98.03	100.72
35			98	0	19240	19292	5	98.19	19123	6	98.03	100.88
36			98	0.5	19292	19404	5	98.19	19203	6	98.03	101.05
37		0.25	90	-0.5	17734.5	17860	5	90.44	17736	5	90.04	100.70
38			90	0	17734.5	17988	5	90.44	17853	5	90.04	100.76
39			90	0.5	17734.5	18116	5	90.44	17970	5	90.04	100.81
40			95	-0.5	19718	19735	5	95.44	19505	6	95.05	101.18
41			95	0	19846	20013	5	95.44	19704	6	95.05	101.57
42			95	0.5	19974	20292	5	95.44	19903	6	95.05	101.95
43			98	-0.5	21691	21693	5	98.30	21299	6	98.05	101.85
44			98	0	22144	22203	5	98.30	21664	6	98.05	102.49
45			98	0.5	22597	22712	5	98.30	21926	7	98.07	103.58
46		0.33	90	-0.5	18528	18636	5	90.60	18419	5	90.04	101.18
47			90	0	18528	18878	5	90.60	18640	6	90.06	101.28
48			90	0.5	18528	19120	5	90.60	18813	6	90.06	101.63
49			95	-0.5	20856.5	20875	5	95.53	20511	6	95.06	101.77
50			95	0	21166	21332	5	95.53	20837	6	95.06	102.38
51			95	0.5	21475.5	21789	5	95.53	21036	7	95.10	103.58
52			98	-0.5	23052.5	22995	4	98.04	22626	6	98.06	101.63
53			98	0	23743	23757	4	98.04	23040	7	98.09	103.11
54			98	0.5	24433.5	24519	4	98.04	23398	8	98.07	104.79
55	1000	0.10	90	-0.5	19375.5	19433	5	90.11	19425	5	90.07	100.04
56			90	0	19375.5	19439	5	90.11	19430	5	90.07	100.05
57			90	0.5	19375.5	19445	5	90.11	19436	5	90.07	100.05
58			95	-0.5	20588	20609	5	95.17	20568	5	95.03	100.20
59			95	0	20588	20645	5	95.17	20601	5	95.03	100.21
60			95	0.5	20588	20681	5	95.17	20634	5	95.03	100.23
61			98	-0.5	21688	21680	5	98.19	21585	5	98.02	100.44
62			98	0	21740	21792	5	98.19	21684	5	98.02	100.50
63			98	0.5	21792	21904	5	98.19	21784	5	98.02	100.55
64		0.25	90	-0.5	20223	20186	4	90.03	20186	4	90.02	100.00
65			90	0	20223	20317	4	90.03	20317	4	90.02	100.00
66			90	0.5	20223	20448	4	90.03	20448	4	90.02	100.00
67			95	-0.5	22097.5	21955	4	95.02	21955	4	95.02	100.00
68			95	0	22197	22235	4	95.02	22235	4	95.02	100.00
69			95	0.5	22296.5	22515	4	95.02	22515	4	95.02	100.00

Exp	k	CV	β	w	$E(TC)_{MILP}$	$E(TC)_{sim}$	NrO		$E(TC)_{MINLP}$	MINLP	Avg β_{MINLP}	$E(TC)_{sim}$
70			98	-0.5	23780	23719	4	98.03	23719	4	98.03	100.00
71			98	0	24216	24227	4	98.03	24227	4	98.03	100.00
72			98	0.5	24652	24735	4	98.03	24735	4	98.03	100.00
73		0.33	90	-0.5	20910	20771	4	90.03	20771	4	90.02	100.00
74			90	0	20910	21016	4	90.03	21016	4	90.02	100.00
75			90	0.5	20910	21261	4	90.03	21261	4	90.02	100.00
76			95	-0.5	22990	22855	4	95.03	22855	4	95.03	100.00
77			95	0	23268	23313	4	95.03	23313	4	95.03	100.00
78			95	0.5	23546	23771	4	95.03	23771	4	95.03	100.00
79			98	-0.5	25052.5	24995	4	98.04	24995	4	98.04	100.00
80			98	0	25743	25757	4	98.04	25757	4	98.04	100.00
81			98	0.5	26433.5	26519	4	98.04	26519	4	98.04	100.00
82 ¹	500	0.25	95	0	20385	20395	6	95.05	20120	7	95.05	101.37
83 ²			95	0	20282	20535	4	95.48	19626	5	95.10	104.63
84			95	0	20037	20277	5	96.19	19495	5	95.06	104.01
85			95	0	19909.5	20137	5	95.67	19486	6	95.07	103.34
86			95	0	20348	20346	4	95.00	19728	6	95.05	103.13

¹ $M = 2$; ² $M = 4$.

The MILP model is considered to be suitable for use in practice if the expected total costs are close to the costs of the optimal solution, and the average fill rate is close to the requirement. The perception of 'close' has to be determined by the producer in practice. Table 6.5 shows that the MILP model performs best with setup cost $k = 500$ or 1000 , when the number of deliveries is limited to 4 or 5 times during the 12 period time horizon. In those cases the expected total costs are less than 5% higher than those of the optimal policy. Interesting is that for setup cost $k = 0$, the MILP policy prescribes for more than half of the experiments not to deliver in each period. Due to aggregation and the fill rate requirement, the total replenishment quantity can be lower when delivering for more periods. If the holding cost is lower than the cost of production of extra items, one delivers for multiple periods. In all experiments but one, the MINLP approach prescribes to deliver in every period. The MINLP approach has lower expected total costs, because it takes the pipeline inventory into account that, apart from period 1, is nonzero, in determining the replenishment quantity. The MILP model determines the replenishment quantity ignoring the availability of pipeline inventory. The evaluated MILP approach goes for certain as the level of the pipeline inventory is highly uncertain due to the long production time. By less replenishments, MILP lowers the total replenishment quantity during the time horizon. With setup cost $k = 1000$ and $CV = 0.25$ or 0.33 , the MILP policy prescribes to deliver every 3 periods. In these cases there is no pipeline inventory and the MILP policy coincides with the optimal policy.

A higher coefficient of variation leads to fewer or the same number of deliveries. The same holds for a higher fill rate. The effect of varying the cost of waste is less clear. The MILP solution is insensitive to varying the tested values of w when $k > 0$. When $k = 0$, varying w

leads to an increase, a decrease or no change in the number of deliveries. For the MINLP solution, varying w leads to an increase or no change in the number of deliveries. Varying the maximum shelf life M confirms the earlier findings with less orders in the MILP solution than optimal and slightly higher costs. The other tested erratic demand patterns have the same number of deliveries than the base demand pattern, while the highly erratic demand pattern performs better on the expected total costs than the erratic variant. The stationary demand requires only four deliveries.

The MILP approach appears especially suitable if the cost structure is such that one does not deliver every period. In that case the pipeline inventory is relatively small and consequently the approach performs well with respect to costs. In case the replenishment cycle is equal to the maximum shelf life, the pipeline inventory is zero, so the MILP approach gives the optimal solution. This analysis is also true for other values of the maximum shelf life.

6.7 Conclusions

We studied the practical problem to determine a production plan for a perishable product with a long lead time and a fixed time horizon under a cycle fill rate constraint. Demand is non-stationary. In case of out-of-stock, demand is lost. Issuing is according to a FIFO policy. We focus on a $Y_t Q_t$ policy, where Y_t denotes in which periods to deliver, and Q_t the corresponding replenishment quantities to deliver. We investigated whether it is possible to construct practical solutions using existing solvers. We considered a single-product – single-echelon SP model, minimizing the expected total costs. To find approximate solutions for the SP model, an MILP model has been formulated. The MILP model is a deterministic approach that generates feasible production plans in less than a second for the performed 86 experiments. With a scenario-based MINLP approach, optimal solutions with respect to a large sample of demand paths are generated as a benchmark for the MILP solutions. The results are data-dependent, but from the performed experiments can be concluded that if the setup cost is low, the MILP model solutions have fill rates higher than required and expected total costs higher than in the optimal solution. Finding a reasonably good approximation for the pipeline inventory in case of non-stationary demand could solve this problem. If the setup cost is higher, such that the replenishment cycle lengths are equal or close to the length of the internal shelf life, the influence of the starting inventory is less. The MILP model generates production plans with fill rates close to the required values and expected total costs are close to optimal. Given the results and the short solver time, the MILP model is suitable for use in practice.

Chapter 7

Conclusion and general discussion

7.1 Introduction

As stated in Chapter 1, in Europe the total food loss and waste is 31% of the initial production from which 6.1% occurs in the food processing, packaging and distribution (HLPE, 2014). Reducing the annual food loss and waste will result in benefits for companies, consumers and the environment in terms of money, volume, energy and sustainability. The aim of this thesis was to contribute to better decision making regarding inventory control in the food supply chain from food producer to supermarket in the trade-off between product-availability and waste. Methods were designed to generate practical order policies using commercial solvers for business rather than custom made solution procedures, where at least the timing of ordering or production is set beforehand.

This chapter is organised as follows. First, in Section 7.2 the conclusions of the five individual research opportunities investigated in the thesis are considered. Section 7.3 focuses on the integrated findings concerning scientific contribution and managerial impact of the research. In Section 7.4 some characteristics and assumptions of the models in the thesis are discussed, followed by directions for future research in Section 7.5.

7.2 Conclusion

The practical inventory control problems discussed in this thesis are characterised by a single perishable product with a fixed lifetime, non-stationary stochastic demand, a single-echelon production/inventory situation in a finite time horizon. Moreover, the product has a service level requirement to ensure a certain product-availability. The decision problems deal with fixed setup or ordering cost, holding cost and disposal cost for wasted items. The age-distribution of the items in stock is considered in specific theoretical Stochastic Programming (SP) problems that deal with the service level requirement as a chance or loss function constraint. The studied problems vary in planning characteristics. The service level is an α -service level constraint on the probability of no stockout, or a fill rate constraint on the proportion of demand delivered directly from stock. The issuing policy is FIFO or combined LIFO – FIFO. When demand exceeds the inventory on hand, demand is either backlogged or alternatively is lost. Lead time of the replenishment is zero, one period or very long.

7.2.1 RO1: MILP approximation generating a YS policy

The first research opportunity was to find a practical approach to determine YS parameter values for a food producer, according to a static-dynamic uncertainty strategy. For this case, the lead time is zero, an α -service level requirement applies and there is FIFO issuing. The research question was whether it is possible to construct practical solutions using commercial solvers for business use rather than custom made solution procedures for this problem.

An MILP model has been formulated for a static-dynamic uncertainty strategy to generate approximate solutions of the formulated SP model. A solution provides a plan specifying simultaneously the periods to produce and the corresponding order-up-to-levels. To meet a certain α -service level, the model considers and corrects for the expected age-distribution of the items in stock, resulting in a waste-compensating replenishment cycle policy. The MILP model can be solved in a fraction of a second. This makes it interesting for practical applications. The generated policy is benchmarked by a flexible $Q(X)$ policy generated by SDP, according to a static uncertainty strategy. For the tested instances, on average the SDP production plan is 1.2% more costly than the MILP production plan.

In situations where the replenishment cycles are either always one period, or always equal to the maximum shelf life, the service levels are exactly fulfilled by the MILP generated policy. The policy has specifically been tested on production plans with replenishment cycles of irregular length. In this situation, the MILP solution provides approximations of the required service level. For these difficult cases, simulation shows that in 96.4% of the periods the service level requirements are met, with an error tolerance of 1%. The performance of the MILP model with respect to the service level is more cumbersome when there are many items of different ages in stock.

7.2.2 RO2: SDP in inventory control: a $Q(X)$ policy

The second research opportunity was to investigate whether Stochastic Dynamic Programming (SDP) is suitable to find policies for inventory systems with service level requirements and non-stationary demand, as is the case for the practical problems under consideration. Inventory control is inherently a multi-stage problem, and therefore SDP seems an appropriate approach to generate policies. The behaviour of SDP for an α -service level constraint and a fill rate constraint was studied. The service level constraints apply per period, being minimal service level constraints. A non-perishable product was considered to avoid the complication of the age-distribution of perishable items in stock. The research question was whether SDP is a suitable method to generate an order policy if a service level constraint applies.

The question was studied by means of a stylized example of an inventory control problem that is relevant for practice. The example has a fixed time horizon with a non-stationary demand distribution. Evaluation of the SDP order policies shows that SDP generates order policies that reach a higher service level than required, because the service level requirement is dealt with in a conditional way depending on the current state of inventory. From all possible starting inventory levels, no matter how small the chance of occurrence, SDP has to meet the service level requirement for the next period. The SDP policy was compared with an order-up-to level policy. The corresponding order-up-to levels were generated by full enumeration meeting the service level requirement, resulting in a policy with lower costs and service levels that are at the end of the replenishment cycles close to the required service level. In general, one can conclude that SDP generates order policies that meet a conditional service level constraint, which does not completely match the

practical service level requirement. Therefore, the resulting policy is not necessarily optimal for a given minimal service level requirement.

7.2.3 RO3: Comparing policies and solution methods

The third research opportunity was to consider the practical problem of a food producer and compare different order policies and different solution methods. The lead time is zero, an α -service level requirement applies and there is FIFO issuing. The research question was in which situations which policy and which solution method is most suitable.

New sample based order policies (control rules) for a SP inventory control problem were developed and compared with existing policies. The inventory control can be handled according to a static-dynamic uncertainty strategy, where a YS policy and a $YQ(X)$ policy were distinguished, with X denoting the inventory at the beginning of the period. Characteristics of a solution to these policies were discussed, followed by a new computational method based on the Smoothed Monte Carlo method with sampled demand, called the $YS_{SMC-MINLP}$ policy and a sample based method to calculate values for the $YQ(X)$ policy. These policies were compared to the YS_{MILP} policy developed in Chapter 2 and a more flexible $Q(X)$ policy generated by SDP, according to a dynamic uncertainty strategy. In most evaluated scenarios, the expected total costs of the policies are very close, and a YS policy gives a cost-efficient and easy to implement solution. In situations of relatively low or high setup cost, MILP generates appropriate parameter values. In situations of intermediate setup cost, where the replenishment cycles are highly varying, the SMC-MINLP parameters might be more suitable. When also the required service level is high and the uncertainty of the demand increases, the age-distribution is important. Furthermore, the flexible policy is only appropriate when setup cost is low, demand is highly uncertain and the required service level is more than 98%.

7.2.4 RO4: YS policy in retail

The fourth research opportunity was to investigate a practical situation in retail, where the reorder days are fixed and order-up-to levels are used, according to a static-dynamic uncertainty strategy. Lead time is one day, an α -service level requirement applies and there is FIFO or combined LIFO – FIFO withdrawal. A lost sales situation applies to retail. There is a weekly demand pattern which is stationary over the weeks. The research question was whether MILP generated parameter values are suitable in case of a lead time of one period and FIFO or combined LIFO – FIFO withdrawal in a rolling horizon application.

Retail order policies were studied for an inventory system of a perishable product with a remaining shelf life on delivery at the store of three days. Customer demand is either FIFO or combined LIFO – FIFO, with a LIFO fraction of 0.4 or 0.6. An SP model of the situation in the retailer practice has been developed. An MILP approximation model is derived, which generates parameter values for a YS order policy. For the practical application of the YS policy, four policies were developed that determine the order quantity based on the YS parameter values found by the MILP model. In the YS_{Σ} policy, the order quantity is

determined considering the total available inventory. This gives lowest cost solutions, which are not always feasible. From the three other policies, $YS_{\gamma\sum t}$ where a fraction of the total available inventory is considered when determining the order quantity, performs best. One can always find a fraction that gives feasible solutions. The calculations take less than a second and can be executed with standard software, which is important for implementation in practice. The developed approach can find a suitable reorder plan according to a static-dynamic uncertainty strategy where the order days are fixed.

7.2.5 RO5: MILP approximation for a YQ policy

The fifth research opportunity was to find a practical approach for a producer of a fresh food product with a long lead time to determine an order policy. Having a long lead time, the inventory levels at the time of delivery are unknown at the moment of deciding the replenishment quantity. A so-called static uncertainty YQ policy was considered, under a cycle fill rate service level requirement and FIFO issuing. In case of out-of-stock, demand is lost. The research question was whether it is possible to generate a production plan for T periods using existing solvers and for which instances the solution is close to the optimal solution.

An MILP model has been formulated to find YQ policy solutions for the specifically formulated SP model. The MILP model is a deterministic approach that generates feasible production plans in less than a second. With a scenario-based MINLP approach, optimal solutions with respect to a large sample of demand paths were generated as a benchmark for the MILP solutions. The results are data-dependent, but from the performed experiments can be concluded that if the setup cost is low, the MILP model solutions have fill rates higher than required and expected total costs higher than the optimal solution. Finding a reasonably good approximation for the pipeline inventory in case of non-stationary demand could solve this problem. If the setup cost is higher, such that the replenishment cycle lengths are equal or close to the length of the internal shelf life, the influence of the starting inventory is less. The MILP model generates production plans with fill rates close to the required values and expected total costs are close to optimal. Given the results and the short solver time, the MILP model is suitable for use in practice.

7.3 Integrated findings

Table 7.1 gives an overview of the research opportunities investigated in this thesis with the corresponding problem characteristics. For every research opportunity, the used methods are listed together with the corresponding order policies.

The starting points for research opportunities RO1, RO3, RO4 and RO5 were practical problems concerning a perishable product with non-stationary demand where the timing of the (production) orders is fixed for planning purposes of a producer or retail organisation. Therefore the aim was to find order policies according to a static-dynamic uncertainty strategy (RO1, RO3 and RO4) and a static uncertainty strategy (RO5). Research opportunity

RO2 investigated an SDP approach to generate order policies according to a dynamic uncertainty strategy as a possible benchmark. The research in this thesis has a scientific contribution as well as managerial impact which will be discussed in Section 7.3.1 and Section 7.3.2.

Table 7.1 Overview of the problem characteristics for each research opportunity and thesis chapter with the used methods and resulting order policies

Problem	RO1 – Ch2	RO2 – Ch3	RO3 – Ch4	RO4 – Ch5	RO5 – Ch6
Characteristics					
Supply chain actor	Producer		Producer	Retailer	Producer
Perishability	Fixed shelf life	∞ life time	Fixed shelf life	Fixed shelf life	Fixed shelf life
Demand distribution	Normal	Uniform/Gamma	Normal	Poisson	Normal
Service level	α -service level	α -serv./ fill rate	α -service level	α -service level	fill rate
Issuing	FIFO	FIFO	FIFO	LIFO – FIFO	FIFO
Excess demand	Backlogging	Lost sales	Backlogging	Lost sales	Lost sales
Lead time	$L = 0$	$L = 0$	$L = 0$	$L = 1$	Long
Method \ Policy					
MILP	YS		YS	YS	YQ
MINLP			YS		YQ
Full enumeration		$Y=1, S$	$YQ(X)$		
SDP	$Q(X)$	$Q(X)$	$Q(X)$		
Sample based approach		$Q(X)$	YS, $YQ(X)$		YQ

7.3.1 Scientific contribution

This thesis studies order policies for a perishable product with non-stationary erratic demand. The fluctuations in demand combined with fixed setup or ordering costs imply that regular production or ordering is probably not optimal. This asks for a strategy to deal with the fluctuations. In inventory literature, non-stationary erratic demand for a perishable product with a fixed lifetime is hardly studied, and consequently also strategies to deal with it are little investigated. To illustrate this, the most recent review of inventory systems of perishable products is due to Bakker et al. (2012). They found 227 papers published between January 2001 and December 2011, from which only 48 assumed a stochastic demand. Half of those papers assumed a fixed lifetime of the product, with a stationary or a price-dependent demand. To link to the situation in practice, the main focus of the thesis was a static-dynamic uncertainty strategy (RO1, RO3 and RO4) and a static uncertainty strategy (RO5), to deal with non-stationary demand. The timing of production or ordering is fixed in advance for planning purposes, while most inventory literature focuses on flexible strategies, which decide every period whether or not to produce or order and how much. This flexible, dynamic uncertainty strategy is also considered in this thesis, in Chapter 3 (RO2), because SDP generated policies are used as a benchmark in Chapters 2 and 4.

Several methods were used to find solutions for the formulated SP models, resulting in order policies for the practical problems. An overview of the used methods is listed in Table

7.1. There are two challenges to face in the SP models: losing inventory due to the perishability of the items causing waste and the service level constraints. The main contributions to the inventory literature are discussed, starting with contributions in the used methods, followed by a general reflection on order policies according to different strategies.

Chapters 2, 4, and 5 focus on a static-dynamic uncertainty strategy, where the order timing is fixed in advance. **MILP approximations** are used to generate parameter values for a YS policy. Considering the literature, MILP has been used before by Tarim and Kingsman (2004) to find exact solutions for inventory control of a non-perishable product with a service level requirement. MILP is a fast way to generate parameter values for a YS order policy with standard software. In Chapter 2, a new MILP formulation has been developed for the case of a perishable product. The main challenge is the formulation of the **set of FIFO constraints** to keep track of the age-distribution in inventory. The investigation showed that using only inventory balance constraints leads to a cost-optimal issuing, which has no clear rule for use in practice. To impose FIFO, an explicit set of inventory constraints with associated logical constraints has been developed, which implements the $\max\{\cdot\}$ operators of the SP model formulation. Because the function $\max\{\cdot\}$ is a convex function, Jensen Inequality (Mood et al., 1974) applies. According to Jensen Inequality the expected waste is underestimated and the expected inventory level of the freshest items is overestimated. Due to the nested function $\max\{\cdot\}$ for the items with an age in-between, Jensen Inequality does not apply and the approximation could be an under- or overestimation. This observation explains why the MILP policy may not meet the service level requirements when replenishment cycle lengths are highly varying and waste occurs during the replenishment cycle.

In Chapter 4, a novel approach in inventory management is described, called the **Smoothed Monte Carlo method**. To generate YS parameter values that are feasible in all periods, this computational method based on MINLP has been developed. For the replenishment cycles where waste will occur during the replenishment cycle, samples of the demand series are used to estimate the value of the order-up-to level S in the service level constraint. This is called the Monte Carlo method. Because the estimator of the service level constraint is piecewise constant, finding the optimal value of S is hard for NLP. The service level approximation as function of the variables S is made practically continuous by a Monte Carlo smoothing approach. For given timing values of Y , NLP finds values of order-up-to levels S close to the optimal ones. A search algorithm finds the best combination of Y and S .

Chapter 6 focuses on a static uncertainty YQ policy under a cycle fill rate service level requirement. The challenge is to **implement a cycle fill rate service level** in an MILP model for different replenishment cycle lengths. The order quantities are determined by solving the loss function of the Normal distribution for the required fill rate. For setup cost that are such that not every period production takes place, this approach appears to be close to the optimal policy obtained by an MINLP approach that generates parameter values approaching optimality with an increasing number of samples.

In Chapter 3, a dynamic uncertainty strategy is considered. The focus is on **the ability of SDP to handle service level constraints** in the context of non-stationary demand inventory control. In some cases, SDP generates RS_t or RS_tS_t policies for a non-perishable product. Evaluation of the SDP order policies shows that SDP generates order policies that overachieve the service level requirements, due to meeting a conditional service level constraint. As illustrated, surprisingly enough there may be order policies that give lower expected costs for a situation of a service level requirement than the SDP generated $Q(X)$ policy. This means, it is not necessarily optimal.

Reflecting on the research in this thesis can be concluded that an optimal policy for the SP model with service level constraints is very hard to find. Fig. 7.1 shows a **classification of order policies** that can be derived from the SP models. The inner circle indicates the static uncertainty strategy where the order timing Y and the order quantity Q is fixed. Relaxing the predetermined order quantity gives a static-dynamic uncertainty strategy as shown in the middle circle of Fig. 7.1. In this strategy, information about the inventory on hand is taken into account. Information about the age-distribution of the perishable inventory improves the solution of a YS policy by compensating for the expected waste during the replenishment cycle. Information about the age-distribution of the perishable inventory is required for a $YQ(X)$ policy where explicitly the age-distribution is used to determine the order quantity. Relaxing also the fixed order timing leads to the dynamic uncertainty strategy in the outer circle with flexible order policies. Most of the inventory literature discusses policies of this type. For RS_t and RS_tS_t policies, information about the age-distribution of the inventory on hand improves the solution. For the $Q(X)$ policy this information is required.

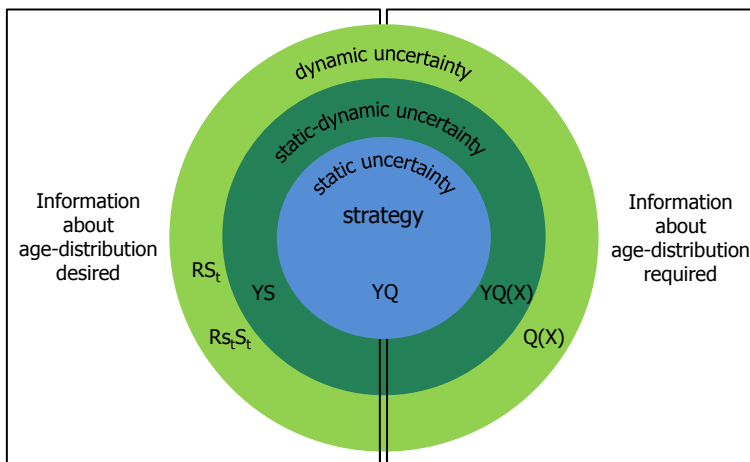


Fig. 7.1 Classification of order policies and required information

Reflecting on the classification of policies from outside in, one can argue that the set of $Q(X)$ policies contains the set of $YQ(X)$ policies, and the set of $YQ(X)$ policies contains the set of

YQ policies, for certain problem characteristics. Note that the YQ policies investigated in this thesis have different problem characteristics compared to the other investigated policies.

Tunc et al. (2013) showed for a non-perishable product with a non-stationary demand that the minimum cost for a static uncertainty strategy are much higher than the minimum cost for a static-dynamic uncertainty strategy. Reducing the so-called quantity-oriented system nervousness comes at a high cost (Tunc et al., 2013), though unavoidable in case of long lead time. However, the cost performance of the static-dynamic uncertainty strategy and the dynamic uncertainty strategy are comparable. The so-called setup-oriented system nervousness can be prevented at a minor cost penalty (Tunc et al., 2013). The results in this thesis give a similar indication for a perishable product.

Regardless of the strategy, a minimum cost policy for the formulated SP models should be of the form $Q(X)$. The optimal order quantity should be flexible and should depend on the age-distribution of the items in stock. Looking at SDP for solving the SP problem, the conditional implementation of the service level requirement for the next period at any state, hinders finding an optimal solution. Some states have a very low probability of occurrence. One would like to assign probabilities to all states in the state space to derive the optimal solution to the problem with a service level constraint. Unfortunately, this seems not possible. The full enumeration approach to generate an optimal $YQ(X)$ policy has the same handicap, as shown in Chapter 4. Studying the SDP solutions, the optimal solution of the SP model will be of the form $Q(X)$, and for a relatively high setup cost it will appear as $s(X)Q(X)$, with a reorder level depending on the inventory composition. The cost-optimal solution will be according to the dynamic uncertainty strategy.

Overall can be concluded that the main contribution of this thesis to the literature of perishable inventory is due to its study of – and design of methods to generate – order policies according to the static-dynamic – and static uncertainty strategy for products with non-stationary demand, where at least the **order timing is fixed in advance**.

7.3.2 Managerial impact

The classification of policies that can be derived from the SP models shown in Fig. 7.1 has also managerial implications as depicted in Fig. 7.2. The static uncertainty strategy fixes the order timing Y and quantity Q in advance. This is **beneficial for planning purposes**, but leaves no flexibility to respond to uncertain outcomes. Tolerating more flexibility in determining the order quantity, but still fixing the order timing according to the static-dynamic uncertainty strategy, is an interesting compromise. It allows detailed planning with some flexibility incorporated to respond to uncertainty. The dynamic uncertainty strategy offers maximal flexibility to be responsive to the uncertainties in demand, but complicates planning.

Perishable inventory literature on order policies according to the static-dynamic uncertainty strategy, with a fixed order timing, is scarce. However, it is known from practice that the

food producer makes a weekly detailed production plan where the decision to produce a certain product or not, is made beforehand. A detailed production plan contains a production schedule with production quantities. In retail, often fixed reorder days of the week exist, to plan warehouse – and store activities. Facing non-stationary demand, the timing of the (production) orders might not be regular, certainly not in retail where ordering every other day is less likely than ordering on fixed days of the week.

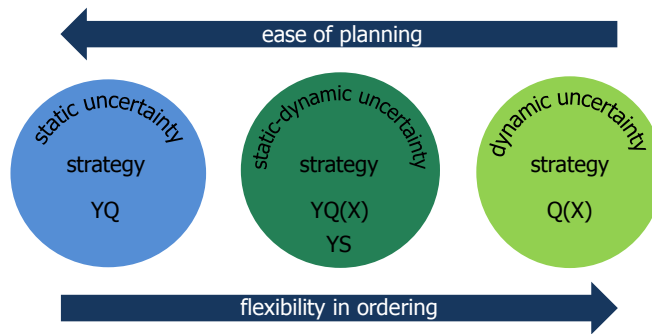


Fig. 7.2 Managerial implication of the choice of policy

This thesis showed that one can generate parameter values for an easy to implement YS policy with standard MILP software in less than a second. In cases where the replenishment cycle length is always equal to 1, or always equal to the maximum shelf life M , MILP generates optimal parameters for a YS policy. We focused specifically on a product with a maximum shelf life of $M = 3$, and situations where the replenishment cycle length varies, because in those situations finding the optimal order policy parameter values is most challenging, due to the occurrence of waste during the replenishment cycle. In many cases, MILP generates parameter values that might be acceptable, being close to the service level requirements. For the other cases, a computational method is developed based on MINLP, called the Smoothed Monte Carlo method, to generate YS parameter values. In Chapter 4 the two YS policies and a $YQ(X)$ policy according to the static-dynamic uncertainty strategy are compared to a dynamic uncertainty strategy $Q(X)$ policy. It is interesting to observe that the expected total costs of all order policies are very close for most scenarios. This corresponds to the findings of Tunc et al. (2013) for a non-perishable product. For a food producer facing a long lead time, the thesis shows that one can generate parameter values for an easy to implement YQ policy. In cases where the replenishment cycle is longer than one period, MILP generates suitable parameter values. For cases where in every period a production run takes place, an MINLP algorithm has been developed to generate parameter values close to optimal. It is the task of management to decide whether it is necessary to use the more accurate MINLP algorithms, which need a longer calculation time, but generate close to optimal parameter values for a YS and a YQ policy.

The policies developed in this thesis according to the static-dynamic uncertainty strategy are applicable in the practice of fresh food producers and in retail organisations, on the level of

the store, as well as on the level of the warehouse. Policies with fixed order timing are **hardly more costly** than the investigated policy according to the dynamic uncertainty strategy **and increase the ease of planning**.

Chapter 1 discussed the topic of food waste. The studied cost functions in the models comprise disposal cost for which the value of the disposal cost in the experiments was varied. The results show that in a static-dynamic uncertainty strategy, increasing the disposal cost can reduce the percentage of waste, maintaining the service level. In setting a (fictitious) value for the disposal cost, one can influence the percentage of waste in determining the order policy, and thus contribute to better decision making in the food supply chain from food producer to supermarket in the trade-off between product-availability and waste.

7.4 Discussion

The studied models incorporate several problem characteristics and underlying assumptions that are worth discussing. This section considers some explicit characteristics and assumptions, as well as more implicit assumptions.

The models considered are finite horizon, periodic review single item – single echelon problems, for a perishable product with a fixed (internal) shelf life M . We mainly consider a **maximum shelf life** of $M = 3$, because this is the most difficult case to solve where on one hand items of different ages can be in stock, and on the other hand often waste during the replenishment cycle occurs. A shorter shelf life is easier to solve as it generates more waste, but it is less likely that waste appears during the replenishment cycle. Therefore the developed model meets the service level requirements better. A longer shelf life culminates less potential waste.

For the perishable product under consideration, a **non-stationary demand** applies. The investigation assumes that for producers the non-stationary demand is erratic due to forecasted promotions, and no bullwhip effect occurs. This assumption is made to highlight that fluctuations in demand occur anyway and have to be dealt with. It is not a limiting assumption to the model. For retailers, a weekly (seasonal) demand pattern which is stationary over the weeks has been considered. This assumption is made to simplify the evaluation of the model. As is shown in the models for the producer, stationarity over the weeks is not a limiting assumption. Also in situations of changing demand due to season or promotion effects, the model can be applied.

In all models, the periods are of equal length. However, depending on the practical situation one can define the **length of the period**, e.g. to a week, a day or part of a day. The length of the period for a fresh food producer will depend on the maximum shelf life of the product. When the shelf life is represented in weeks like in cheese production, a period will be a week. When the shelf life is represented in days, it is likely that it is more convenient to define a period being a day. In retail stores, it is custom to order on fixed days of the week.

A product as packaged cheese has a shelf life of a few weeks. Therefore standard ordering without considering the perishability will not cause problems. However, for products that are days-fresh (Van Donselaar et al., 2006), like e.g. fresh cut lettuce or fresh mussels, it is important to consider the perishability and the period in the model should be a day.

The models of Chapter 2 and 4 that were formulated for a food producer assume a period is a week and **lead time** is zero, so there is immediate production. A lead time of zero is often used in inventory literature, because it simplifies calculations. In reality, production takes place during the week. Therefore the lead time is not zero, but can be smaller than one week, e.g. several days. We can consider this as a lead time of zero. For practice it is important that the production weeks are fixed in advance, with an estimation of the production quantity. The actual production quantity is determined at the beginning of the week, based on the order-up-to level and the inventory-on-hand. It is even possible to adapt the order-up-to level to recently updated forecasts. In such a situation a lead time of zero should not be a limitation. However, if a period is a day, one should consider adding lead time to the model, as is done in Chapter 5.

In the models formulated from a producer perspective, we apply a **time horizon** of 12 periods, being 12 weeks. This horizon is based on considerations in practice where a rough production planning contains at most 8 weeks. For a longer time horizon, the forecast of demand will be less reliable. To avoid end-of-horizon effects in the model solutions, a time horizon of 12 weeks is used. Combined with a maximum shelf life of 3 periods, the model generates at least 4 replenishment cycles. At least at every start of a new replenishment cycle the model calculations will be executed again, with updated demand forecasts.

The model formulated from a retailer perspective (Chapter 5) assumes stationary demand over the weeks. That makes a time horizon of 7 days the logical choice. However, in case of non-stationary demand over the weeks, a longer time horizon can be chosen, e.g. 14 days. The MILP solver can easily handle the extra binary variables.

In all research opportunities, **minimal service level requirements** apply. In considering product-availability, service level requirements reflect the situation in practice well, as discussed in Section 1.2.3. In the static-dynamic uncertainty strategy one can find feasible solutions meeting the service level requirements. Nevertheless, service levels cause a challenge if one wants to find the optimal solution. In inventory control models, SDP often generates optimal solutions. As shown in this study, it seems very hard to capture the service level constraints in SDP such that SDP generates an optimal solution. One can choose to use a shortage cost to consider product-availability instead of a service level constraint. Then SDP can generate an optimal solution. Van Houtum and Zijm (2000) showed a relationship between shortage cost and a mean service level in case of stationary demand. For non-stationary demand and minimal service level constraints this relation is not straightforward. The challenge is that there is no one-to-one relationship between a minimal service level and shortage cost in case of non-stationary demand.

The studied models are **single item** models. However, in practice one needs to consider many products, depending on the supply chain actor and the production stage in case of a food producer. At the packaging stage, one raw material can be used to produce multiple end products. The number of stock keeping units depends highly on the type of food producer. For all products, inventory control is necessary. Moreover, in the production planning of multiple products, the products will compete on resource capacity.

In retail, a supermarket may have to consider about 10,000 days-fresh stock keeping units. For example, a supermarket of the Dutch retail chain Jumbo in general has an assortment of about 32,000 stock keeping units (Jumbo, 2016). About one third of the assortment concerns days-fresh products (personal communication, 2016). That means that the calculation time to determine the order quantity has to be less than a second to be suitable for practical use in retail. This holds for the MILP generated parameter values in this thesis, but not for the MINLP and SDP algorithms.

As mentioned in Chapter 1, (Silver et al., 1998) distinguish **system control costs** as part of the relevant costs in inventory management. System control costs are not included in the developed models, but they cannot be disregarded. Calculation time can be considered as system control cost. Depending on the supply chain actor and the amount and costs of different products under consideration, it is important to include the calculation time in the selection of an appropriate order policy. Furthermore, information on the age-distribution of the inventory on hand can be considered as system control cost. In retail stores, the information about the age-distribution is not always electronically available. However, having this information will generate better reorder decisions.

An implicit assumption in the models in this thesis is that there is **ample capacity** to produce or deliver the ordered products. In retail practice, the volume of the combined orders may be fairly constant, so one can argue that the capacity of the trucks used for delivery can be adapted to the total volume of the orders, such that the assumption will not limit the applicability of the model. For a food producer, the assumption of ample capacity can generate solutions that exceed the capacity.

The models in this thesis are developed according to the static-dynamic uncertainty strategy, for situations where a **fixed setup or ordering cost** applies. For a food producer setting up a production run costs time and money and therefore it might not be (economically) feasible to run a production for a slow moving product in every planning period. For a fast moving product where setup cost is low relative to the holding cost, production may be scheduled every period such that order policies according to a dynamic uncertainty strategy are suitable. In retailer practice the same reasoning applies. For large supermarkets with an every day delivery, order policies according to a dynamic uncertainty strategy are suitable. However, small stores that are delivered on fixed days of the week, reflecting a static-dynamic uncertainty strategy, need an order policy that considers replenishment cycles of different lengths. When fixed ordering cost is not considered for the

supermarket, limiting the amount of orders per week instead of incurring a fixed ordering cost gives the same results.

7.5 Directions for future research

The integrated findings and the discussion lead to a few directions for future research. The focus of this thesis is on order policies according to a static-dynamic – or a static uncertainty strategy. However, for a fast moving product and for benchmark reasons it would be interesting to find optimal order policies according to a dynamic uncertainty strategy for a perishable fixed lifetime product with a non-stationary demand and a service level requirement. A minimum cost policy for the formulated SP models should be of the form $Q(X)$. The optimal order quantity should depend on the age-distribution of the items in stock. Up to now, we were not able to find optimal order policies with SDP for the described problem, due to the conditional implementation of the service level requirement. The question is whether it is possible to **find optimal order policies** according to a dynamic uncertainty strategy with a service level requirement and which method would be suitable.

In a situation of long lead time, a static uncertainty strategy is required. Adaptation of the order quantity just before realisation of demand is not possible, so the timing and production quantity have to be determined at the beginning of the planning horizon. When the setup cost is relatively low, such that a production is started in almost every period, there will be a considerable pipeline inventory causing the MILP model solutions have fill rates higher than required and expected total costs higher than the optimal solution. The question is how to find a reasonably good **approximation for the pipeline inventory** in case of non-stationary demand in a fast way.

It would be interesting for future research to investigate whether the inventory models for a food producer can be extended to **multiple products**, including capacity constraints.

In retail stores, the **information about the age-distribution** of the items on shelf is not always electronically available. However, having this information will generate better reorder decisions. A cost-benefit analysis of keeping records of the best-before dates of sold products could be carried out.

The models in this thesis are designed for practical inventory control problems and evaluated in a simulation with data based on practical situations. However, it would be interesting to test the models with **company data**.

Finally, decisions on order timing and order quantity downstream in the supply chain influence the performance of the actors upstream in the supply chain. Also when the retail organisation fixes the order timing for the supermarket, the performance on cost and waste of the supermarket is influenced. The models studied in this thesis are single-echelon models. It would be interesting for future research to extend the inventory models to **multiple echelons** to model these effects and support coordination in the supply chain.

References

Summary

Acknowledgements

References

- Alcoba, A.G., Hendrix, E.M.T., García, I., Ortega, G., Pauls-Worm, K.G.J., Haijema, R., 2015. On Computing Order Quantities for Perishable Inventory Control with Non-stationary Demand, in: Gervasi, O., Murgante, B., Misra, S., Gavrilova, M.L., Rocha, A.M.A.C., Torre, C., Taniar, D., Apduhan, B.O. (Eds.), *Computational Science and Its Applications - ICCSA 2015*. Springer International Publishing, pp. 429-444.
- Axsäter, S., 2006. *Inventory control*. Springer Science & Business Media.
- Bakker, M., Riezebos, J., Teunter, R.H., 2012. Review of inventory systems with deterioration since 2001. *Eur. J. Oper. Res.* 221, 275-284.
- Bellman, Richard, and E. Stanley Lee. 1978. "Functional equations in dynamic programming." *Review of. aequationes mathematicae* 17 (1):1-18. doi: 10.1007/bf01818535.
- Bijvank, M., Vis, I.F.A., 2011. Lost-sales inventory theory: A review. *Eur. J. Oper. Res.* 215, 1-13.
- Bijvank, Marco, and Iris F. A. Vis. 2012. "Lost-sales inventory systems with a service level criterion." *Review of. European Journal of Operational Research* 220 (3):610-8. doi: 10.1016/j.ejor.2012.02.013.
- Bookbinder, J.H., Tan, J.Y., 1988. Strategies for the probabilistic lot-sizing problem with service-level constraints. *Manage. Sci.* 34, 1096-1108.
- Broekmeulen, R., van Donselaar, K.H., 2009. A heuristic to manage perishable inventory with batch ordering, positive lead-times, and time-varying demand. *Comput. Oper. Res.* 36, 3013-3018.
- Burgin, T. A. 1975. "The Gamma Distribution and Inventory Control." *Review of. Operational Research Quarterly (1970-1977)* 26 (3):507-25. doi: 10.2307/3008211.
- Chao, X., Gong, X., Shi, C., Zhang, H., 2015. Approximation Algorithms for Perishable Inventory Systems. *Operations Research* 63, 585-601.
- Chen, F.Y., Krass, D., 2001. Inventory models with minimal service level constraints. *Eur. J. Oper. Res.* 134, 120-140.
- Chopra, S., Meindl, P., 2010. *Supply Chain Management: Strategy, Planning, and Operation*, Fourth Edition ed. Pearson Education, Inc., New Jersey.
- FAO, 2011. *Global food losses and food waste - Extent, causes and prevention*, Rome.
- FAO, 2015. *Global initiative on food loss and waste reduction*.
- Fries, B.E., 1975. Optimal Ordering Policy for a Perishable Commodity with Fixed Lifetime. *Operations Research* 23, 46-61.
- Goldberg, D.A., Katz, D.A., Lu, Y., Sharma, M., Squillante, M.S., 2014. Asymptotic Optimality of Constant-Order Policies for Lost Sales Inventory Models with Large Lead Times. *arXiv preprint arXiv:1211.4063v2* (2014).
- Goyal, S.K., Giri, B.C., 2001. Recent trends in modeling of deteriorating inventory. *Eur. J. Oper. Res.* 134, 1-16.
- Graves, S.C., Willems, S.P., 2000. Optimizing Strategic Safety Stock Placement in Supply Chains. *Manufacturing & Service Operations Management* 2, 68-83.
- Graves, S.C., Willems, S.P., 2008. Strategic inventory placement in supply chains: Nonstationary demand. *Manufacturing & Service Operations Management* 10, 278-287.

- Haijema, R., 2013. A new class of stock-level dependent ordering policies for perishables with a short maximum shelf life. *Int. J. Prod. Econ.* doi:10.1016/j.ijpe.2011.05.021.
- Haijema, R., Minner, S., 2015. Stock-level dependent ordering of perishables: A comparison of hybrid base-stock and constant order policies. *Int. J. Prod. Econ.*
- Haijema, R., van der Wal, J., van Dijk, N.M., 2007. Blood platelet production: Optimization by dynamic programming and simulation. *Comput. Oper. Res.* 34, 760-779.
- Haijema, R., van Dijk, N., van der Wal, J., Sibinga, C.S., 2009. Blood platelet production with breaks: optimization by SDP and simulation. *Int. J. Prod. Econ.* 121, 464-473.
- Hendrix, E.M.T., Haijema, R., Rossi, R., Pauls-Worm, K.G.J., 2012. On Solving a Stochastic Programming Model for Perishable Inventory Control, in: Murgante, B., Gervasi, O., Misra, S., Nedjah, N., Rocha, A.C., Taniar, D., Apduhan, B. (Eds.), *Computational Science and Its Applications – ICCSA 2012*. Springer Berlin Heidelberg, pp. 45-56.
- Hendrix, E.M.T., Olieman, N.J., 2008. The smoothed Monte Carlo method in robustness optimization. *Optimization Methods and Software* 23, 717-729.
- Hendrix, E.M.T., Pauls-Worm, K.G.J., Rossi, R., Alcoba, A.G., Haijema, R., 2015. A Sample-Based Method for Perishable Good Inventory Control with a Service Level Constraint, in: Corman, F., Voß, S., Negenborn, R.R. (Eds.), *Computational Logistics*. Springer International Publishing, pp. 526-540.
- HLPE, 2014. Food losses and waste in the context of sustainable food systems. A report by the High Level Panel of Experts on Food Security and Nutrition of the Committee on World Food Security, Rome.
- INFORMS, 2016. <https://www.informs.org/About-INFORMS/What-is-Operations-Research> (25-4-2016).
- Jacod, J., Protter, P., 2004. *Probability Essentials*, 2 ed. Springer, Berlin.
- Jumbo, 2016. (<http://www.jumbo.com/content/feiten-en-cijfers/>) (20-3-2016).
- Kall, P., Wallace, S.W., 1994. *Stochastic Programming*. John Wiley & Sons, Chichester.
- Karaesmen, I.Z., Scheller-Wolf, A., Deniz, B., 2011. Managing Perishable and Aging Inventories: Review and Future Research Directions, in: Kempf, K.G., Keskinocak, P., Uzsoy, R. (Eds.), *Planning Production and Inventories in the Extended Enterprise*. Springer, pp. 393-436.
- Karlin, S., 1960a. Dynamic Inventory Policy with Varying Stochastic Demands. *Manage. Sci.* 6, 231-258.
- Karlin, S.C.F.p.d.D., 1960b. Optimal Policy for Dynamic Inventory Process with Stochastic Demands Subject to Seasonal Variations. *Journal of the Society for Industrial and Applied Mathematics* 8, 611-629.
- King, A.J., Wallace, S.W., 2012. *Modeling with Stochastic Programming*. Springer Science + Business Media, New York.
- Lee, H.L., Padmanabhan, V., Whang, S.J., 1997. Information distortion in a supply chain: The bullwhip effect. *Manage. Sci.* 43, 546-558.
- Lyman Ott, R., Longnecker, M.T., 2001. *An Introduction to Statistical Methods and Data Analysis* 5ed. Duxbury, Pacific Grove.
- Minner, S., Transchel, S., 2010. Periodic review inventory-control for perishable products under service-level constraints. *OR Spectrum* 32, 979-996.
- Mood, A.M., Graybill, F.A., Boes, D.C., 1974. *Introduction to the theory of statistics* Third Edition ed. McGraw-Hill.

- Morton, T.E., 1969. Bounds on the Solution of the Lagged Optimal Inventory Equation with no Demand Backlogging and Proportional Costs. *SIAM Review* 11, 572-596.
- Morton, T.E., 1978. The Nonstationary Infinite Horizon Inventory Problem. *Manage. Sci.* 24, 1474-1482.
- Morton, T.E., Pentico, D.W., 1995. The Finite Horizon Nonstationary Stochastic Inventory Problem: Near-Myopic Bounds, Heuristics, Testing. *Manage. Sci.* 41, 334-343.
- Nahmias, S., 1975. Optimal Ordering Policies for Perishable Inventory-II. *Operations* 23, 735-749.
- Nahmias, S., 1982. Perishable inventory-theory - a review. *Operations Research* 30, 680-708.
- Neale, J.J., Willems, S.P., 2009. Managing Inventory in Supply Chains with Nonstationary Demand. *Interfaces* 39, 388-399.
- Pauls-Worm, K.G.J., Hendrix, E.M.T., 2015. SDP in Inventory Control: Non-stationary Demand and Service Level Constraints, in: Gervasi, O., Murgante, B., Misra, S., Gavrilova, M.L., Rocha, A.M.A.C., Torre, C., Taniar, D., Apduhan, B.O. (Eds.), *Computational Science and Its Applications -- ICCSA 2015*. Springer International Publishing, pp. 397-412.
- Pauls-Worm, K.G.J., Hendrix, E.M.T., Alcoba, A.G., Haijema, R., 2015. Order quantities for perishable inventory control with non-stationary demand and a fill rate constraint. *Int. J. Prod. Econ.*
- Pauls-Worm, K.G.J., Hendrix, E.M.T., Haijema, R., van der Vorst, J.G.A.J., 2014. An MILP approximation for ordering perishable products with non-stationary demand and service level constraints. *Int. J. Prod. Econ.* 157, 133-146.
- Pujawan, I.N., Silver, E.A., 2008. Augmenting the lot sizing order quantity when demand is probabilistic. *Eur. J. Oper. Res.* 188, 705-722.
- Rijkema, W.A., Hendrix, E.M.T., Rossi, R., van der Vorst, J.G.A.J., 2013. Application of stochastic programming to reduce uncertainty in quality-based supply planning of slaughterhouses. *Annals of Operations Research*, 1-12.
- Rossi, R., 2013a. On Service Level Measures in Stochastic Inventory Control, in: Natalia, B., Kirill, C., Alexandre, D., Vladimir, L. (Eds.), *Proceedings of the IFAC Conference on Manufacturing Modeling, Management and Control, MIM 2013*, June 19-21. International Federation of Automatic Control, Manufacturing Modelling, Management, and Control, Saint Petersburg, Russia, pp. 1991-1996.
- Rossi, R., 2013b. Periodic review for a perishable item under non stationary stochastic demand, pp. 2021-2026.
- Rossi, R., Tarim, S.A., Hnich, B., Prestwich, S., 2008. A global chance-constraint for stochastic inventory systems under service level constraints. *Constraints* 13, 490-517.
- Rossi, R., Tarim, S.A., Bollapragada, R., 2011a. Constraint-Based Local Search for Inventory Control Under Stochastic Demand and Lead Time. *INFORMS Journal on Computing*.
- Rossi, R., Tarim, S.A., Hnich, B., Prestwich, S., 2010. Computing the non-stationary replenishment cycle inventory policy under stochastic supplier lead-times. *Int. J. Prod. Econ.* 127, 180-189.
- Rossi, R., Tarim, S.A., Hnich, B., Prestwich, S., 2011b. A state space augmentation algorithm for the replenishment cycle inventory policy. *Int. J. Prod. Econ.* 133, 377-384.

- Silver, E.A., Pyke, D.F., Peterson, R., 1998. Inventory Management and Production Planning and Scheduling, Third Edition ed. John Wiley & Sons.
- Sobel, Matthew J., and Rachel Q. Zhang. 2001. "Inventory Policies for Systems with Stochastic and Deterministic Demand." *Operations Research* 49 (1):157-62. doi: 10.1287/opre.49.1.157.11197.
- Tarim, S.A., Dogru, M.K., Özen, U., Rossi, R., 2011. An efficient computational method for a stochastic dynamic lot-sizing problem under service-level constraints. *Eur. J. Oper. Res.* 215, 563-571.
- Tarim, S.A., Kingsman, B.G., 2004. The stochastic dynamic production/inventory lot-sizing problem with service-level constraints. *Int. J. Prod. Econ.* 88, 105-119.
- Tekin, E., Gurler, U., Berk, E., 2001. Age-based vs. stock level control policies for a perishable inventory system. *Eur. J. Oper. Res.* 134, 309-329.
- Tempelmeier, H., 2007. On the stochastic uncapacitated dynamic single-item lotsizing problem with service level constraints. *Eur. J. Oper. Res.* 181, 184-194.
- Tempelmeier, H., 2011. A column generation heuristic for dynamic capacitated lot sizing with random demand under a fill rate constraint. *Omega* 39, 627-633.
- Tempelmeier, H., Herpers, S., 2011. Dynamic uncapacitated lot sizing with random demand under a fillrate constraint. *Eur. J. Oper. Res.* 212, 497-507.
- Tunc, H., Kilic, O.A., Tarim, S.A., Eksioglu, B., 2011. The cost of using stationary inventory policies when demand is non-stationary. *Omega* 39, 410-415.
- Tunc, H., Kilic, O.A., Tarim, S.A., Eksioglu, B., 2013. A simple approach for assessing the cost of system nervousness. *Int. J. Prod. Econ.* 141, 619-625.
- Van der Vorst, J.G.A.J., Beulens, A.J.M., van Beek, P., 2000. Modelling and simulating multi-echelon food systems. *Eur. J. Oper. Res.* 122, 354-366.
- Van Dijk, Diana, Eligius M. T. Hendrix, Rene Haijema, Rolf A. Groeneveld, and Ekko C. van Ierland. 2014. "On solving a bi-level stochastic dynamic programming model for analyzing fisheries policies: Fishermen behavior and optimal fish quota." *Review of. Ecological Modelling* 272 (0):68-75. doi: 10.1016/j.ecolmodel.2013.09.015.
- Van Donselaar, K., de Kok, T., Rutten, W., 1996. Two replenishment strategies for the lost sales inventory model: A comparison. *Int. J. Prod. Econ.* 46-47, 285-295.
- Van Donselaar, K., van Woensel, T., Broekmeulen, R., Fransoo, J., 2006. Inventory control of perishables in supermarkets. *Int. J. Prod. Econ.* 104, 462-472.
- Van Houtum, G. J., and W. H. M. Zijm. 2000. "On the relation between cost and service models for general inventory systems." *Review of. Statistica Neerlandica* 54 (2):127-47. doi: 10.1111/1467-9574.00132.
- Van Zyl, G.J.J. 1964. "Inventory control for perishable commodities." In *Unpublished PhD dissertation*. University of North Carolina.
- Veinott Jr, A.F., 1963. Optimal stockage policies with non-stationary stochastic demands. *Multistage Inventory Models and Techniques*, Stanford University Press, Stanford, CA, 85-115.
- Veinott Jr., A.F., 1965. Optimal Policy for a Multi-Product, Dynamic, Nonstationary Inventory Problem. *Manage. Sci.* 12, 206-222.
- Wagner, H. M., and T. M. Whitin. 1958. "Dynamic version of the economic lot size model." *Review of. Management Science* 5 (1):89-96.

Zipkin, P., 1989. Critical Number Policies for Inventory Models with Periodic Data. *Manage. Sci.* 35, 71-80.

Summary

Globally, around one-third of the edible parts of perishable food products is wasted every year. Adequate logistics management of the food supply chain is of importance. Inventory control of processed fresh food with a best-before or use-by date deals with the questions how much to order and how often, and how to review the inventory position. For perishable products, there exists a trade-off between desired product-availability and waste. A high order quantity or frequent ordering may result in high product-availability but also waste. A small order quantity or less frequent ordering results in lower inventory levels, but may cause out-of-stock. For the inventory control of a perishable product the age-distribution of the items should be considered. This is influenced by order picking at a food producer or warehouse, or consumer behaviour at a supermarket, picking First In First Out (FIFO), where first the oldest items are used, or Last In First Out (LIFO), where first the freshest items of the product are used. This thesis investigates periodic review order policies for a food producer and a retailer for a perishable product with a fixed lifetime and a non-stationary demand. For a food producer as well as a retailer, it is not always economically feasible to reorder in every period. For planning purposes it is desirable to set the timing of the orders in advance.

The aim of this thesis is to contribute to better decision making regarding inventory control in the food supply chain from food producer to supermarket, in the trade-off between product-availability and waste. Methods are designed to generate practical order policies using commercial solvers for business rather than custom made solution procedures, where at least the timing of ordering or production is set beforehand. The practical inventory control problems discussed in this thesis are characterised by a perishable product with a fixed lifetime, non-stationary stochastic demand, a single-echelon production/inventory situation in a finite time horizon. The product has a service level requirement to ensure a certain product-availability. The decision problems deal with fixed setup or ordering cost, holding cost and disposal cost for wasted items. The age-distribution of the items in stock is considered in specific theoretical Stochastic Programming (SP) problems that deal with the service level requirement as a chance constraint or fill rate constraint. The studied problems vary in planning characteristics.

Chapter 2 studies the practical production planning problem of a food producer. In meeting the uncertain demand, the food producer uses a FIFO issuing policy. The food producer aims at meeting a certain α -service level at lowest cost and needs an order policy specifying beforehand the periods to produce and the corresponding order-up-to levels. The formulated SP model assumes zero lead time and backlogging of shortages. A Mixed Integer Linear Programming (MILP) model is constructed that provides an approximate optimal plan for such a policy. The order-up-to level for each period is corrected for the expected waste by explicitly considering for every period the expected age-distribution of the products in stock. The viability of the approach is illustrated by numerical experiments. Simulation

shows that in 96.4% of the periods the service level requirements are met with an error tolerance of 1%.

Chapter 3 investigates whether Stochastic Dynamic Programming (SDP) is suitable to find policies for inventory systems with service level requirements and non-stationary demand. Inventory control is inherently a multi-stage problem, and therefore SDP seems an appropriate approach to generate policies. Small instances are constructed to study the behaviour for an α -service level constraint and a fill rate constraint per period, being minimal service level constraints. To get a clear view on the system dynamics of the service level constraints, a non-perishable product is considered, to avoid the complication of the age-distribution of perishable items in stock. Evaluation of the SDP order policies shows that SDP generates order policies that reach a higher service level than required, because the service level requirement is dealt with in a conditional way depending on the current state of inventory. The SDP policy is compared with an order-up-to level policy. The order-up-to levels were generated by full enumeration meeting the service level requirement, resulting in a policy with lower costs and service levels that are at the end of the replenishment cycle close to the required service level. In general, one can conclude that SDP generates order policies that meet a conditional service level constraint not necessarily matching the original service level requirement. Given a minimal expected service level requirement, the resulting policy is not necessarily optimal.

Chapter 4 returns to inventory control of a perishable product. The food producer uses a FIFO issuing policy and wants to meet an α -service level requirement at lowest cost. Various inventory policies can be defined to handle the inventory control problem, from easy to implement to more complex decision rules. Three different policies are studied: a) the production timing is fixed in advance combined with an order-up-to level, b) the production timing is fixed in advance and the production quantity takes the age-distribution into account and c) every period is decided to produce or not, and how much, depending on the age-distribution of the items in stock. The study in Chapter 2 used an MILP approximation to derive values for policy a). In Chapter 4 a computational method is used, based on the so-called Smoothed Monte Carlo method with sampled demand to optimize values for such a system. The resulting Mixed Integer Non-Linear Programming (MINLP) approach uses enumeration, bounding and iterative nonlinear optimisation. A sample based approach is presented, to determine the order quantities for policy b). Flexible policy c) is derived by SDP. All policies are compared. In most tested scenarios the expected total costs of the policies are very close and the MILP values for policy a) provide a cost efficient and easy to implement solution. In situations where the cost structure is such that the replenishment cycle lengths are highly varying, waste will occur during the replenishment cycle and the parameters generated by MINLP may be more suitable. Only in situations of a high service level and increasing uncertainty of demand, the age-distribution is important. Furthermore, the flexible policy is only appropriate when setup cost is low, demand is highly uncertain and the required service level is more than 98%.

Summary

Retail order policies are studied in Chapter 5, for a practical situation where stores may be able to order every day, or on fixed days of the week. Focus is on a perishable product which has a remaining shelf life on delivery at the store of three days and a weekly (seasonal) demand pattern which is stationary over the weeks. The stores have a target α -service level. The customer demand may be FIFO or a LIFO – FIFO combination. Four different policies to determine the order quantity are studied. The base is an order policy where the reorder days are fixed and order-up-to levels are used, which vary per day of the week. The parameter values are generated by an MILP approximation of the specifically formulated SP model. Numerical experiments compare the effectiveness of the policies to determine the order quantities with respect to costs and reached service levels. The first policy determines the order quantity considering the total available inventory. This gives lowest cost solutions, which are not always feasible. From the three other studied policies, the policy where a fraction of the total available inventory is considered when determining the order quantity, performs best. One can always find a fraction that gives feasible solutions.

In Chapter 6 the practical decision problem of fresh food production with a long production lead time is studied, to decide every period (e.g. week) how many items to produce. When a batch is ready for use, its items have a fixed shelf life, after which the items become waste in the sense that they cannot be sold anymore. The focus is on cases where a so-called cycle fill rate service level requirement applies. The generation of a production plan that fixes the timing and quantity of the production for a finite time horizon is investigated. To minimise waste, one uses a FIFO issuing policy. In case of out-of-stock, sales are lost. To find approximate solutions for the formulated SP model for this case, an MILP model is developed. With a scenario-based MINLP approach, optimal solutions are generated for a large sample of demand paths as a benchmark for the MILP solutions. The MILP model is suitable for practical use if the setup cost is such that the replenishment cycles in the production plan are close to or of the same length as the maximum shelf life. In those cases, the expected total costs are close to the costs of the optimal solution and the average fill rate is close to the required one.

To conclude, this thesis studies order policies for a perishable product with non-stationary demand. The fluctuations in demand combined with fixed setup or ordering costs imply that regular production or ordering is probably not optimal. This situation requires a strategy to deal with the fluctuations. In inventory literature, non-stationary demand for a perishable product with a fixed lifetime is hardly studied, and consequently also strategies to deal with it are little investigated. In this thesis, methods have been designed to generate practical order policies using commercial solvers, where at least the timing of ordering or production is set beforehand. The practical order policies have been benchmarked with custom made solution procedures. The developed practical policies are applicable in the practice of fresh food producers and in retail organisations, on the level of the store, as well as on the level of the distribution centre. An interesting conclusion is that the policies with fixed order timing are hardly more costly than the investigated policy with flexible order timing. The fixed order timing increases the ease of planning.

The studied cost functions in the models comprise disposal cost for which the value of the disposal cost in the experiments was varied. The results show that in the models with fixed order timing and corresponding order-up-to levels, increasing the disposal cost can reduce the percentage of waste, maintaining the service level. In setting a (fictitious) value for the disposal cost, one can steer the percentage of waste in determining the order policy, and thus contribute to better decision making in the food supply chain from food producer to supermarket in the trade-off between product-availability and waste.

Acknowledgements

During my thesis project I learned a lot about the topic of my thesis, about doing research, about my colleagues and about myself. I had a lot of support from my colleagues from Operations Research and Logistics and from Information Technology. I thank you all for your support, for the cooperation in research and education and the good times we had and still have during work. I really appreciate it to be part of a nice group of colleagues.

A few colleagues I would like to thank personally. First, I would like to thank Jack, my promotor. When you became chair of the Operations Research and Logistics group you brought new energy and new research projects. This inspired me to extend my employment from 40% to 60% and to start working on my PhD thesis. You gave me the opportunity to choose my own topic. Without you I would have never started this thesis.

Secondly, I acknowledge the co-authors of the papers in this thesis, Eligius, René, Roberto, Alejandro and Jack, for the cooperation and their contribution.

Furthermore, I would like to thank Joke and Frits, for working together, for your encouragement and for supporting me also during the defence. It is good to know that there will be colleagues by my side who work even longer at our group than I do. 😊 Natasja, Ilona and Jeanette, I offer you many thanks for all your help and attention, always, but especially in the last phase of my thesis project.

I thank Joke, Maarten and Jacqueline (in order of appearance) for our lunch walks. I value the walks and talks very much as they offer fresh air during the work, both literally and metaphorically. I learned a lot from all of you in different ways. Thank you for that. I hope to continue these walks as long as we work together.

Last but not least, I would like to thank Eligius. When I started the thesis, you moved to Spain. You were less involved at the start of my thesis, but at some point you came in again and never went away. You flew in many times from Spain, and I really liked working with you. Your enthusiasm helped me a lot, and I enjoyed our sparring sessions and cooperation. You were always supportive and optimistic. Without you this thesis would not (yet) have been finished. Thank you!

Karin Pauls

Arnhem, August 2016



Name of the learning activity	Department/Institute	Year	ECTS*
A) Project related competences			
Writing Research Proposal	WU	2008	6
Revenue Management	LNMB	2008	4
Stochastic Programming with Applications in Logistics	Aarhus University	2013	5
B) General research related competences			
Introduction course	MG3S	2008	1.5
Techniques for writing and presenting scientific papers	WGS	2008	1.2
<i>Computing replenishment cycle policy parameters for a perishable item with non-stationary stochastic demand under service level constraints</i>	EURO international conference, Lisbon	2010	1
<i>Replenishment cycle policy parameters for perishable items under service level constraints</i>	LNMB, Lunteren	2011	1
<i>Production planning of a perishable product with lead time and non-stationary demand</i>	ISIR international conference, Budapest	2012	1
<i>Production planning of a perishable food product with non-stationary demand in a rolling planning horizon</i>	OR2013, international conference, Rotterdam	2013	1
<i>Order policies for a perishable product with non-stationary demand under a fill rate constraint</i>	ISIR international conference, Budapest	2014	1
<i>Fixed reorder days for a perishable product in retail</i>	ISIR international conference, Budapest	2016	1
C) Career related competences/personal development			
Basiscursus didactiek	OLU-LUW	1999	2.4
Afstudeervak Organiseren en Begeleiden	OWU-WU	2005	0.6
Procesbegeleiders AMC	OWU-WU	2004	0.6
Developing/coordinating the course ORL-30806	WU	2003-2016	4
Culturele diversiteit in het onderwijs	OWU-WU	2005	0.6
Mobilising your Scientific Network	WGS	2015	1
Total			32.9

*One credit according to ECTS is on average equivalent to 28 hours of study load

Colophon

This research has been supported and partly funded by

1. TI Food and Nutrition (project RE002), The Netherlands (Chapters 2 and 6)
2. The European Community's Seventh Framework Programme (FP7) under grant agreement no 244994 (project VEGi-TRADE) (Chapter 2)
3. The Spanish state (project TIN2012-37483-C03-01), in part financed by the European Regional Development Fund (ERDF) (Chapters 2, 3 and 6)
4. Junta de Andalucía (P11-TIC-7176), in part financed by the European Regional Development Fund (ERDF) (Chapters 2, 3 and 6)
5. The Spanish Ministry (TIN2015-66680-C2-2-R), in part financed by the European Regional Development Fund (ERDF) (Chapter 4 and 5)

Cover photo by Aleksandr Bryliaev/Shutterstock

Printed by Digiforce / Proefschriftmaken.nl, Vianen