On Inventory Management
for Products with a Stochastic Selling Season

Inaugural Dissertation
to Obtain the Academic Degree
of a Doctor in Business Administration
at the University of Mannheim

submitted by

Luca Biscaro
Mannheim

Dean: *Joachim Lutz*

Referent: *Prof. Dr. Moritz Fleischmann*

Correferent: *Prof. Dr. Raik Stolletz*

Day of oral examination: *22 June 2022*

Acknowledgments

This thesis marks the end of a five-year long journey. The help, support and love of many people I have had the honor of meeting and having alongside me in the joyful and challenging parts of this journey are what made it possible for me to complete it. To all these people I owe a debt of gratitude.

First of all, I would like to thank Moritz Fleischmann, my PhD advisor. I am especially grateful to him for teaching me not to forget the big picture irrespective of the stage of the research, and to start simple when facing new and complex tasks. His guidance, advice, exceptional intuition and enthusiasm kept this project on track and ensured its successful completion. In addition, he created a family-like atmosphere at the chair, which made my time as a PhD student that much easier and more enjoyable.

Furthermore, I am thankful for the openness, enthusiasm and professionalism of all those involved in the real-world case study that we were able to conduct at an agrochemical company. This collaboration inspired and shaped this thesis.

I would also like to thank all my fellow PhD candidates and colleagues in the Operations Management area of the University of Mannheim. The lunch breaks with Matteo and Johannes, the laughs and talks with Amir and my office mate Hossein, and the kindness and friendliness of everybody else put a smile on my face even in the most challenging days.

In the same vein, I am grateful to my friends in Italy for their long-standing and long-distance friendship. Every time I come back home, they make me feel like I have never left, and there is nothing I enjoy more in life than spending time with them.

Finally, I feel extremely lucky to have had my family beside me in this journey. Most of all, I thank my parents for their unconditional love and support, not only during this time, but in my entire life. I dedicate this thesis to them.

Contents

List of Figures

List of Tables

1 Introduction

1.1 Research topic and contributions

The inventory management problem for products with a seasonal and uncertain demand has long been a subject of interest in the operations research literature, the newsvendor model (Arrow et al., 1955) being not only the central model of this specific stream of operations research literature, but also a core model of operations research in general. This interest is due to the high practical relevance of the problem, since many products in a common make-to-stock setting have a demand that is both seasonal and subject to some degree of uncertainty. In this thesis, we study inventory management in a seasonal demand setting. Inspired by the experience acquired conducting a real-world case study in the agrochemical industry, our focus is on addressing the dynamics of the selling season.

Using the terminology of Mitchell (1927), demand seasonality can be caused by either climate or convention. Naturally, climatic seasons are a driver of seasonality for the demand for many products: for example, the demand for skis and ski apparel is heavily concentrated in the late autumn and winter, whereas the demand for grilling equipment increases in spring and summer. In comparison, the demand for these products in the remainder of the year is negligible. Conventions, which can have many origins, namely religious holidays, fashions and business practices, can also be the source of seasonal variations in demand: for example, the demand for Christmas trees and decorations peaks in December, whereas the demand for costumes is highest before festivities such as Carnival and Halloween. Even with limited or no uncertainty, demand seasonality is an important inventory driver for many companies. Indeed, the need for seasonal inventory could in principle be eliminated by building a production capacity large enough to meet the peak season's demand just in time. However, such a large capacity would require a substantial investment and be inefficient in the low season when it is not needed. Therefore, companies build up inventory during the low season periods to meet the peak season's demand and maximize their profits.

Similarly, demand uncertainty, a ubiquitous characteristic of demand, is an important inventory driver in an make-to-stock environment, even in the absence of seasonality. To maximize expected profits or to reach service-level targets, companies create buffers against demand uncertainty with safety stock. The optimal safety stock quantity is the one which strikes the optimal balance between expected underage and overage costs, i.e. the costs of producing less and more than the demand, respectively, or the quantity which meets the required service level at minimum cost.

When a product's demand is both seasonal and uncertain, the inventory problem faced by the company becomes significantly more intricate, especially when the demand's seasonality is driven by climatic seasons. In these cases, it is important to explicitly consider and manage two types of demand uncertainty: quantity and timing (Schlapp and Fleischmann, 2020). Not only the total volume of demand during the selling season is uncertain, but also the start and duration of the season. In addition, customers' demand patterns during the season can also significantly fluctuate between seasons. These two uncertainty types imply that the firm needs to make two related inventory decisions – a quantity and timing decision, respectively. Similar to the quantity decision, the timing decision involves finding the optimal balance between two types of expected costs – those of making the product available for sale before and after the start of the selling season. Making the product available too early causes the firm to incur costs of holding inventory until the product is demanded. In contrast, making the product available too late, to avoid these costs, causes the firm to lose out on demand and profits in the early part of the season. Moreover, the quantity and timing decisions are clearly interconnected, and cannot be made independently of each another. Making the product available early allows to potentially capture the entire season's demand, justifying holding a large inventory, whereas with a later inventory timing the expected demand in the remainder of the season is lower, incentivizing holding a smaller inventory. As a result, when the demand for a product is seasonal and subject to quantity and timing uncertainty, the firm must build up inventory to meet the peak season's demand without knowing its exact volume and, especially when the seasonality is caused by climatic seasons, timing. Because of limited production capacities, an investment in seasonal and safety inventory must be made in advance, despite the firm's information on the quantity and timing of demand possibly improving as the peak season approaches, under certain circumstances, thereby diminishing the uncertainty.

Although there is substantial literature on inventory management for products with seasonal and uncertain demand, the focus has been on the quantity uncertainty of demand and the related inventory quantity decision, while the timing uncertainty of demand and the inventory timing decision have not received much attention. The newsvendor model, the starting point of this stream of operations research literature, assumes that the season is instantaneous (has a length of zero) and that its start is known with certainty. A number of extensions to this central model do consider either the timing uncertainty of demand (e.g. Hadley, 1962, Nahmias, 1977) or the inventory timing decision (e.g. Choi et al., 2004 and Ravindran, 1972), however, they do not consider both elements simultaneously. To the best of our knowledge, Schlapp and Fleischmann (2020) are the only authors to consider both uncertainties and both decisions simultaneously. These authors are the first to analytically characterize a firm's optimal inventory timing and quantity decisions, and the relationship between them, when the seasonal demand for the firm's product is subject to both quantity and timing uncertainty. Moreover, in the literature, a method to efficiently manage this specific inventory problem in practical settings is also lacking, as are case studies. Real-world settings are more complex than the ones considered in the – mostly – analytical studies found in the literature, because of factors such as multiple products sharing the same production resources, and more complex production processes, both of which give rise to additional inventory trade-offs. To capture these complexities, a multi-period discrete time view appears most appropriate.

In this thesis, we study the inventory problem of a firm selling products with seasonal and uncertain demand, with the properties of the season itself being stochastic. We contribute to the literature in several ways. First, we build on the findings of Schlapp and Fleischmann (2020) to show how the problem's parameters influence the firm's inventory strategy – the combination of quantity and timing – and how considering incorrectly or neglecting altogether the timing uncertainty affects the firm's profits. Second, to capture the aforementioned complexities found in practice, we go beyond a stylized newsvendor-like modeling approach. Specifically, we develop a reformulation of the scenario approximation of the stochastic capacitated lot-sizing model (SCLSP), which is a planning model widely used in practice. The stochastic model adopts a discrete multi-period view, and allows to make production and inventory decisions for multiple products sharing the same production resources' capacity and to manage three common and important inventory drivers – demand seasonality, demand uncertainty and economies of scale. Our reformulation is able to significantly simplify the demand distribution estimation and scenario-generation procedures in the case of demand dependencies over time and across products, which naturally characterize the seasonal inventory problem setting considered in this dissertation. Indeed, when demand is uncertain, seasonality implies that the demand in different time intervals is linked by the shape of the season (the demand pattern), thus leading to demand autocorrelation. Moreover, cross-correlation also naturally occurs in this setting, because some products may be demanded in the same (climatic) season or in distinct consecutive seasons. Third, we adapt this newly proposed methodology and apply it in analyzing a real-world inventory management problem faced by an agrochemical producer. Agrochemicals, or crop-protection products, are perfect examples of products with a highly seasonal and uncertain demand, both in terms of quantity and timing. Indeed, these products' demand follows the highly seasonal agricultural and pests' life cycles, whose exact timing, which depends on weather conditions, is difficult to predict. The uncertainty of the demand for these products is exacerbated by its dependence on crops' prices, as well as the availability and price of substitute agrochemicals in the market. Moreover, the complexity, costs and long lead times of the production processes of agrochemicals make it necessary for inventory decisions to be made well ahead of the start of the selling season. Therefore, correctly managing the inventory quantity and timing decisions is crucial for the success of agrochemical producers.

1.2 Structure of the thesis

The remainder of this thesis is structured as follows:

In Chapter 2, we provide an in-depth description of the inventory management problem that agrochemical producers confront. This chapter serves two purposes. First, it provides a practical, concrete example of the inventory problem analyzed throughout this thesis and thereby serves as a motivation. Second, it introduces the setting of the real-world case study presented in Chapter 5.

In Chapter 3, we build on the results of Schlapp and Fleischmann (2020), complementing their analysis in two ways. First, we show how, and to what extent, the problem's parameters affect the optimal inventory quantity and timing decisions, and the company's profits. Second, we quantify the benefits of considering both the timing and quantity uncertainties of demand when setting the inventory strategy. This is in comparison to using a common naïve inventory policy which neglects timing uncertainty and the role of inventory timing in managing the economic trade-offs arising from the combination of both aforementioned types of demand uncertainties. To achieve this, we vary the values of a set of parameters of the problem within a specific practically relevant range and conduct the resulting full factorial numerical study. We solve each of the total $5,283$ instances of the problem using both a naïve approach based on the classical newsvendor model and a (near) optimal approach, which consists of a grid-search procedure to solve a scenario approximation of the stochastic inventory problem. We use analysis of variance (ANOVA) techniques to analyze the results. First, we show that the effect of varying the problem's parameters related to the stochastic properties of the product's selling season, within the specified ranges, is larger on the optimal inventory timing decision than on the quantity component of the inventory strategy of the firm. Furthermore, the opposite is true for the cost-related parameters of the problem. We then show that the naïve inventory policy, which neglects the timing uncertainty of demand, leads to substantially lower expected profits than the optimal policy. This is especially true when the timing-related properties of the season are subject to a higher degree of uncertainty, the inventory costs are higher and the season lasts longer on average.

In Chapter 4, we continue the study of the inventory problem central to this thesis by focusing on its practical application. We make four important modifications to the more stylized problem setting of Chapter 3. First, we consider the case of a company with limited production capacity that produces multiple products, each with a potentially different stochastic selling season. Second, we assume a more complex production process, in which setups are necessary before the production of a product can start. Third, we consider more complex demand patterns, modeling, to this end, the problem in discrete time, a common approach employed in practice, as opposed to continuous time. Fourth, we allow backlogs and consider the problem of determining the production and inventory strategy that minimizes expected costs while meeting a desired service-level target, instead of maximizing expected profits. To solve this problem, we develop a reformulation of the scenario approximation of the SCLSP, using a method that we named the cumulative demand scenario (CDS) approach, which is based on building scenarios for the cumulative demand of the products, as opposed to the standard approach used in the literature of building scenarios that are complete paths in so-called scenario fans or trees. We show that this approach can significantly simplify the demand distribution estimation and sampling procedures in the case of demand dependencies over time and across products. As previously explained, this is especially useful in settings where demand is seasonal and the properties of the season are stochastic, because demand dependencies are naturally present and complex. Moreover, we conduct a numerical study under the more general assumption of dynamic stochastic demand which shows that, compared to previously proposed scenario-approximation methods, the newly proposed approach is more computationally efficient and more robust to changes in the parameters of the problem, such as the coefficient of variation of demand.

In Chapter 5, we study the real-world mid-term inventory problem for crop protection products faced by an agrochemical company. As it is common in the agrochemical industry, the demand for these products is highly seasonal and uncertain, both in terms of quantity and timing. We model the inventory problem as a stochastic general lot-sizing and scheduling problem (SGLSP), because this enables us to capture the main characteristics of the company's complex production processes, and, based on the nature of the demand for the products, we use the newly developed CDS approach, presented in Chapter 4, to approximate and solve it. Using this model and solution approach, we provide the agrochemical company with valuable insights into their inventory planning problem. First, we show that demand uncertainty is the company's dominant inventory driver. This emphasizes the importance of investing resources into improving the management of demand uncertainty, for example by improving demand forecast accuracy. Second, through a simulation using actual data from three past seasons, we show that our stochastic model and solution approach outperforms the planning method currently used by the company and clearly identify the source of the latter's shortcomings. Third, we provide the company with multiple tools to support and possibly improve their planning process in the future.

Finally, in Chapter 6, we conclude by providing a summary of the work done, the results achieved, and the opportunities for future research.

2 Inventory management at agrochemical producers: a motivating example

To illustrate the relevance of the problem of managing the inventory of products with a stochastic seasonal demand, in this chapter we present and discuss in detail the inventory management challenges faced by agrochemical producers. These companies face a highly seasonal and uncertain demand for most of their products and, therefore, making accurate inventory quantity and timing decisions is crucial to their economic performance. The discussion in this chapter also serves as an introduction to Chapter 5, where we present the results of a mid-term inventory problem case study conducted at an agrochemical company.

The remainder of this chapter is organized as follows. In Section 2.1 we present a general introduction to agrochemicals and the agrochemical industry. In Section 2.2 we describe the supply chain (SC) of agrochemical producers and discuss the inventory planning problem they face. To conclude, in Section 2.3 we present a literature review on SC planning in the agrochemical industry.

2.1 Agrochemicals and the agrochemical industry

Agrochemicals, also known as crop protection products (CPPs), or pesticides, are chemical products used in the agricultural sector to protect crops from harmful organisms (pests) or diseases (European Commission, 2020), and mainly include herbicides, fungicides and insecticides. The main benefit to applying CPPs is to preserve the crops on which they are applied and thereby improve productivity. CPPs are manufactured and sold in the form of formulations. The formulation process combines the active ingredients (AI) of the CPP, i.e. the chemicals which actively control the pests, with inactive ingredients, which make the product safer, more effective and easier to use. The primary end-users of CPPs are typically micro and small agricultural enterprises, which directly apply the products to their crops. The right product to apply, as well as the timing and quantity of its application depend on which stage of their life-cycles the crops and/or the pests are in, the severity and spread of the infestation, and the weather and soil conditions before, during and after application. These are the factors that influence the efficacy of CPPs (Bouma, 2003).

When used inappropriately, agrochemicals can negatively impact the environment, as well as the health of the users of the products and of the consumers of the food derived from the crops on which they are applied. Therefore, the agrochemical industry is highly regulated. Accordingly, products are required to undergo a complex and strict approval procedure before they are allowed to be distributed in any country. This procedure, which differs by region, is usually lengthy and must be regularly repeated. For example, in the

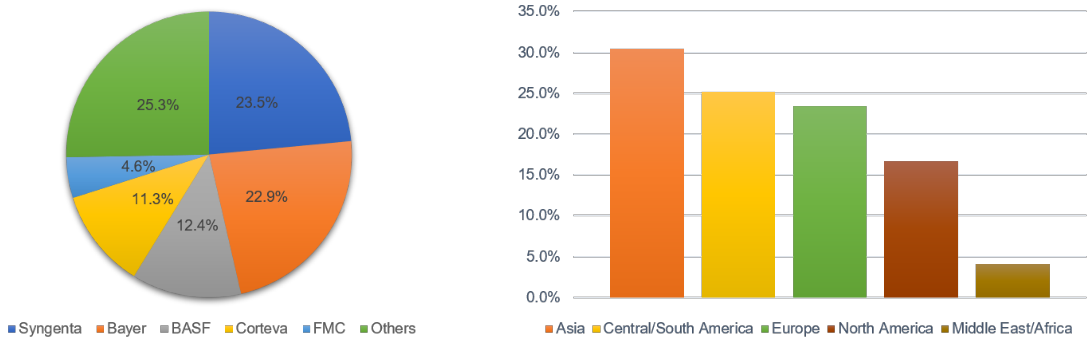

(a) Market share by company

Source: Mooney (2018)

(b) Market share by region

Source: AgbioInvestor as cited in Phillips (2020)

Figure 1: The agrochemical industry in 2018

European Union, the approval process of a new active ingredient lasts on average three years and seven months, and if successful, the active ingredient is approved for use for a maximum of 15 years (European Commission, 2020). New CPPs, i.e. formulations, must also be authorized by each individual member state in which they are to be distributed, with the authorization requiring periodic renewal as well. Moreover, the requirements for approval and authorization are constantly changing and have globally become more stringent in recent decades. Currently, approvals might be reconsidered before their expiration and further restrictions can be placed on the use of the products after they have been approved (see European Commission, 2016).

The progressively complex regulatory environment has increased the average time needed for a new product to enter the market after being discovered, which in 2014 took on average 11 years (Sparks and Lorsbach, 2017). This is an increase of about 36% over the course of the previous two decades. Coupled with a low probability of success in the discovery and development process of agrochemical compounds, which is currently estimated as one in 160 thousand, this has caused the monetary investment necessary to discover and develop a new agrochemical to rise significantly. Recent estimates suggest that it now costs on average around $286 million to discover and develop a new agrochemical (Sparks and Lorsbach, 2017). Considering the necessity to develop new products due to the growing resistance of pests to agrochemicals, these increasingly high research and development (R&D) costs are one of the factors which contributed to a consolidation of the agrochemical industry in recent decades, with companies investing 7% to 10% of their sales in R&D yearly (Nishimoto, 2019; Sparks and Lorsbach, 2017). At present, the five leading agrochemical producers in the global market are Syngenta, Bayer, BASF, Corteva and FMC and in 2018 their total market share was nearly 75%. These producers are big multinational companies characterized by a high degree of vertical integration (Fritz and Hausen, 2009), and, given that they operate and sell globally, they have large and complex SC networks. In Figures 1a and 1b we report the share of the CPP market by company and region, respectively, as of 2018.

In addition, regulations affect agrochemical companies by imposing strict rules on the production, transportation and storage of agrochemicals, complicating SC planning activities and operations. For example, the country of origin of the active ingredients of a CPP might decide where the product can be sold and the amount of sales taxes applied to it (Bassett and Gardner, 2013). Moreover, some countries might have special formulation and packaging requirements for CPPs (Fritz and Hausen, 2009).

2.2 Agrochemical companies' supply chains and planning challenges

In Figure 2 we present a simplified scheme of the typical SC of CPPs. The first part of the network, up to warehousing, is a typical agrochemical company's SC, whereas the second part, shown for completeness, is the remaining part of the SC until the product reaches its end-users. The general SC structure shown and, therefore, many of the planning problems presented later in this section are shared by all the major players in the agrochemical industry. Indeed, as previously stated, the market is dominated by large multinational companies from the process industry which share many similarities and are active in the same markets.

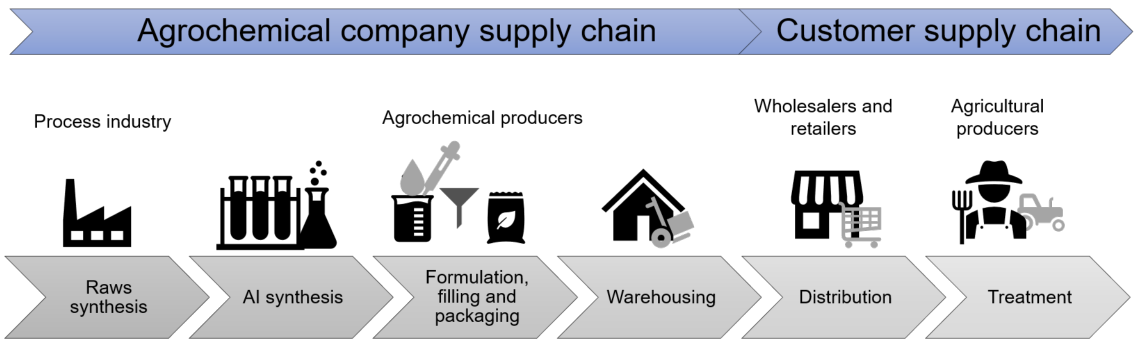

Figure 2: Typical supply chain of crop protection products. Adapted from Fritz and Hausen (2009)

The first step in the chain is raws synthesis and refers to producing raw materials that are necessary to produce AIs. The suppliers of raw materials are generally companies belonging to the process industry. The raw materials are then synthesized into AIs in the next stage of the SC process. AI production is typically performed in a single and sometimes dedicated plant. The production processes in these plants are carried out either continuously or in campaigns (Fritz and Hausen, 2009), i.e. long periods of time in which multiple batches are produced successively to cover the demand for a long time horizon (even an entire season), which are then followed by a downtime period until the next campaign starts. Moreover, AI synthesis can take between a few months and up to two years (Shah, 2004). Therefore, this step of the production process of agrochemicals has an important impact on the SC, as it causes a high replenishment time for the next stages and creates substantial inventory. It is also clear that AI synthesis is a "push" activity, i.e. it is based on demand forecasts rather than actual customers' orders, because AI production decisions (both timing and quantity) must be taken long before the selling season. The AI production planning problem, especially in the case of campaign production, therefore, resembles the classic newsvendor problem. Once AIs are produced, they are transported to formulation plants, where the final agrochemical products are produced. Formulation plants usually produce a wide range of products, which in some cases could belong to different product families, i.e. herbicides, fungicides or insecticides. The formulation process can be complex and is usually performed in batches, with potentially long changeover operations in order to avoid contamination issues between different products produced successively on the same resource, depending on the production sequence (Shah, 2004). Although more flexible than AI production, formulation is still largely a push activity, due to demand seasonality, time-consuming changeover operations and limited capacity. This means that the production schedule and quantities must be decided in advance of the selling season when the demand for the product is not fully known. Flexibility in this production stage is also limited to avoid planning nervousness. Packaging and labeling, which conclude the production process, can be performed either at formulation sites or dedicated sites (Sousa et al., 2008; Elimam, 1995). In general, formulation, packaging and labeling are organized in either a centralized or decentralized structure (Fritz and Hausen, 2009). Finally, packaged formulations are sent to regional or local warehouses, where they are stored before being sold. Agrochemical companies' direct customers are not the end users of CPPs, but typically medium-sized distributors (Fritz and Hausen, 2009). These distributors then manage the distribution and sale of the products to the ultimate users, i.e. agricultural enterprises.

We note that pharmaceutical companies have similar SCs to that described in Figure 2, due to the similarities between the production processes of pharmaceuticals and agrochemicals. In addition, considering the multiple similarities between the industries, for example the strong impact of regulations and the significant R&D expenditures, pharmaceutical and agrochemical companies face many similar planning problems (see, e.g., Shah, 2004).

One of the major challenges faced by agrochemical companies is production and inventory planning. This is due to the strong demand seasonality, the high degree of demand uncertainty and the complex nature of SC processes, as previously illustrated, that require most activities to be performed based on demand forecasts. The demand for agrochemicals is discontinuous and seasonal (Fritz and Hausen, 2009), as it depends on the growing season of the crops on which they are applied. Due to the limited production capacity and the costs of matching supply with demand, agrochemicals must be produced in the off-season, when there is almost no demand for them, in order to build up the inventory needed to meet the demand during the main selling season. To complicate matters further, different agrochemicals have different selling seasons, in terms of timing, and the selling season of a given agrochemical differs from region to region. Overall, seasonality leads to substantial inventory quantities and costs in the SC. Apart from being strongly seasonal, demand is also difficult to predict, both in terms of timing and quantity. Indeed, the right timing and quantity of CPPs to apply to a certain crop are exceedingly difficult to predict, because the damage caused by pests depends on the complex interaction of two or more living organisms (Rosenzweig et al., 2001). Many CPPs are designed to be applied at a certain stage of the growth cycle of the crop and/or of the life cycle of the pests, and the exact timing of these stages can vary across seasons because of weather conditions. In general, the appearance, spread and severity of pest infestations depend on factors such as temperature, precipitation, humidity, radiation and many other difficult-to-predict weather conditions (Rosenzweig et al., 2001).

For some products, special weather conditions might shift the yearly peak demand period by as much as one or two months (Bloemen and Maes, 1992). Weather conditions also affect the productivity of crops, as well as the performance of CPPs (Bouma, 2003). The uncertainty of demand for agrochemicals is exacerbated by other factors. For example, farmers' demand depends also on the value of the crops, as the investment in CPPs is justified only if the value of the additional yield obtained by its application is larger than its price, and this value is uncertain itself (Böcker and Finger, 2017; Desai, 1970). Moreover, the demand for a company's products is influenced by the availability and prices of the products' substitutes from competitors (Fritz and Hausen, 2009), which are challenging to predict. As a result of the strong uncertainty in demand, the amount of safety stocks required in agrochemical companies' SCs to meet customers' high service level requirements can be considerable.

This discussion emphasizes that agrochemical companies face challenging production and inventory planning decisions for each CPP in their product portfolios. They must decide when and how much AIs and formulated products to manufacture for the season(s) ahead based on insecure forecasts of the start and duration of the selling season and of the demand quantity. Making the product available early allows a company to serve the demand of customers in case of an early start of the season, thus enabling it to capture a larger share of the season's total demand. However, early availability also results in a higher risk of incurring unnecessary inventory costs before the start of the season. Because an early timing decision increases the potential demand that a firm can satisfy, the optimal quantity decision is clearly dependent on the timing decision and vice versa, thus making it necessary to manage both these decisions simultaneously. Moreover, the inventory problem of agrochemical companies is complicated by the fact that the timing and quantity decisions for different products are also closely interconnected, because of the shared limited production capacity, the potentially long and costly sequence-dependent setup operations and the interrelation between the products' selling seasons.

2.3 Literature review of SC planning in the agrochemical industry

Despite their variety and complexity, the SC planning challenges, including inventory management, faced by agrochemical companies have not received much academic attention. Some of these challenges are shared by companies in other industries, for example the pharmaceutical industry, and therefore approaches developed to solve similar SC planning problems in other settings could also be useful and relevant to SC planning for agrochemical companies. However, other challenges are unique to the agrochemical industry and call for tailored solution approaches. In this section, we present a review of the operations research literature on the specific topic of SC planning in the agrochemical industry. This is the literature that the case study of Chapter 5 contributes to. We organize the literature review into long-term (strategic), mid-term (tactical) and short-term (operational) planning problems according to the typical classification of planning problems (see e.g. Fleischmann et al., 2015).

2.3.1 Long-term planning literature

Most of the literature on SC planning in the agrochemical industry focuses on strategic problems, specifically supply chain network design and capacity planning.

Sousa et al. (2008) present a case study conducted at a multinational agrochemical company whose goal is to simultaneously solve a strategic and tactical planning problem. The problems are to redesign the global formulation and distribution SC network of a set of herbicides derived from the same two AIs, and to define an annual cyclic production and distribution schedule for the SC, respectively, with the objective of maximizing profits. Uncertainty of demand is not considered. The need for a cyclic plan is due to the yearly seasonality of demand for the formulated products considered. To solve these problems, the authors develop a hierarchical planning approach, because of the size of the problem. In the upper level of the hierarchy, a mathematical model is used to optimize the configuration of the SC and the aggregate production and distribution plan with a planning horizon of one year, which corresponds to one seasonal demand cycle. In the lower level, multiple detailed production planning models with a time horizon of one month are used to check the feasibility and accuracy of the decisions of the aggregate upper-level model. Feedback from the lower level is provided to the upper level in an iterative fashion to improve the design and planning decisions. The feedback from the detailed model to the aggregate high-level one is given in the form of modified capacity, utilization and demand coverage constraints. Importantly, one of the additional details of the production processes captured by the lower-level model, compared to the higher-level model, is the need for changeover operations. The authors show that the developed iterative hierarchical approach provides more realistic and reliable results than a classical hierarchical approach in which feedback of the detailed planning model is not incorporated into the aggregate model.

Bassett and Gardner (2013) present a large-scale MILP developed to optimize the design of the SC of a group of products having a common AI at Dow AgroSciences. Because of the characteristic seasonality of demand for agrochemicals, the production schedule of the facilities and the distribution schedule within the SC are also modeled and optimized, despite the strategic nature of the project. The main contribution of the paper is the development of a modeling approach that can track the origin of the products sold (down to the raw material level), which, due to the highly regulated nature of the agrochemical industry, determines whether the product can be sold in a given market or not and the amount of duties and taxes applied to it. This ensures that the resulting design and schedule of the SC network is both feasible and profit-optimal. The authors present a case study conducted at the company, which shows that correctly tracking the origin of the products has a significant impact on the solution and corresponding profits.

Bassett (2018) discusses different SC design planning problems faced by Dow AgroSciences. The author presents the different software they typically use to solve them and the input required. Then, he provides a high-level description of three exemplary planning problems. Although the mathematical models used to solve them are not provided, the main features of the problems and the procedure followed to obtain the solutions are presented. The first problem analyzed is that of designing the SC of a newly-developed AI. Even at this strategic level, the author states that seasonality of demand is the main issue to address in the model. Seasonality causes the optimal production schedule to be highly intermittent, whereas usually the company prefers to operate plants in a campaign mode; however, campaigns cause an increase in the inventory within the SC. This makes it necessary to make a decision on the operating mode of the plants before the SC network design model is solved. The second problem analyzed is that of finding the optimal SC design and schedule for multiple products sold in the Asian market. Here the author focuses on the importance and challenges of correctly calculating the duties and taxes which apply to the products sold. The third and last problem presented is a rail fleet optimization problem for the transportation of materials between three plants in the U.S. The author emphasizes again the importance of considering the seasonal nature of demand, and, additionally, the need of considering the uncertainty of demand, which is done using a sensitivity analysis.

Schnelle (2000) studies the smaller-scale problem of designing AI plants at Dow AgroSciences in the 1990's. The author uses a hierarchical approach inspired by Subrahmanyam et al. (1994). The upper-level model is an aggregate design model whose aim is to find the design of the plant, including the number of production resources to buy and the timing of their installation, and the production schedule of the equipment that optimize the expected net present value. Uncertainty in both prices and demand for the AIs is considered and modeled with the use of scenarios. To ensure that the resulting optimal production schedule is feasible, lower-level scheduling subproblems are solved. The solution method used is a heuristic that considers the important changeover times of the resources which produce multiple products, among other aspects. After presenting the approach, the author discusses its application in a real-world project at Dow AgroSciences.

Another important topic studied in the long-term planning literature is that of capacity planning. Liu and Papageorgiou (2013) consider the capacity, distribution and production planning problem of the global SC of an agrochemical company. The main goal is to provide support for capacity expansion decisions at a set of formulation plants, given that demand is forecasted to be larger than the current production capacity of the plants. Demand is assumed to be deterministic. The problem is modeled as an MILP with multiple (conflicting) objectives: minimize total costs, which include raw material, production, transportation and inventory costs, and duties; minimize the total flow time, which measures the responsiveness of the SC; minimize the lost sales, which measure the customer service level of the company. Two methods for solving multiobjective problems are considered: the ε-constraint and the lexicographic minimax methods. A new approach is developed to transform the lexicographic minimax problem into a minimization problem. Finally, the authors present the results of a numerical analysis in which they compare two different expansion strategies using the developed models.

Liu and Papageorgiou (2018) study the problem of fair distribution of profits among members of a three-echelon agrochemical SC for one AI. The authors first present a MILP for production, distribution and capacity planning that optimizes the total profits of the SC, and subsequently develop different approaches to additionally optimize transfer prices, i.e. prices charged for the sale of final and intermediate products among members of the SC, considering two fairness criteria and different bargaining powers of the members. The developed solution approaches are tested on two examples, including a real-world case study, which show that they can achieve a fairer profit distribution than the one obtained by simply optimizing total profits.

Apart from supply chain design and capacity planning, the long-term planning problem of new product development (NPD) has also attracted significant academic attention. Maravelias and Grossmann (2001) study the problem of simultaneous planning of testing activities in NPD and designing and planning of production facilities at agrochemical and pharmaceutical companies. The authors present a two-stage stochastic programming model whose main goal is to find the most promising NPD projects in the R&D portfolio of the company to pursue. The uncertainty of the outcome of different testing tasks is captured through scenarios. Because the net present value of different investment opportunities is closely linked to the time to market, the design of production facilities where the new products will be manufactured must be decided well before it becomes clear whether the product is successful or not. In addition, the right investment in manufacturing assets depends on the production schedule of the new and existing products which should be followed to serve the uncertain future demand. As a result, the model developed by the authors optimizes not only the NPD decisions, but also the decisions of opening new plants or expanding existing ones, and the production schedule of these facilities. Due to the size of the resulting MILP, the authors propose a heuristic solution approach and test its quality in three illustrative examples.

Bassett (2000) studies another aspect of the NPD problem. The author presents a mathematical program developed for Dow AgroSciences to optimally assigns employees to projects in the current R&D pipeline of the company. The objective of the model is to best utilize the employees' expertise and minimize the use of external contractors in the R&D process. To obtain solutions of good quality in acceptable time for what-if analysis, the author further develops a heuristic solution approach. Given that it is outside of the scope of this thesis, we do not survey the NPD literature further; for additional reference, we refer the interested reader to the studies by Subramanian et al. (2000) and Schmidt and Grossmann (1996).

2.3.2 Mid-term planning literature

The mid-term planning literature focuses mostly on production and inventory planning. Bloemen and Maes (1992) present a linear programming (LP) model developed to support mid-term production planning at a Monsanto's plant producing herbicides, in which both AI synthesis and formulation are performed. To account for the seasonality of the demand and the limited production capacity, the planning horizon of the problem is 12-18 months. The objective of the model is to find the optimal balance between inventory costs and the costs of alternative production resources to serve the highly seasonal demand, considering production and storage constraints, as well as the availability of containers used to store or transport products to the demand zones. Demand uncertainty is not considered and the model is kept deterministic. The level of detail of the production processes is limited to keep the model linear, but their multiple-stage nature is reflected. Moreover, the model also considers transportation flows between the plant under consideration and other plants of the company, with a particular focus on the decision of renting containers used to store or transport the products. After presenting the model formulation, the authors discuss the successful implementation of the tool at the company.

Similarly to Bloemen and Maes (1992), Elimam (1995) presents a case study on the application of an LP model to support mid-term production and procurement planning at a US pesticide manufacturing plant. AI and raw materials are procured from vendors, and the plant only performs the formulation step of the production process. The planning horizon considered is 18 months, and the goal of the model is to find the cost-minimizing production plan to meet the highly seasonal future demand. To consider the high variability of demand, safety stock targets for every product and period are exogenously defined. Finally, the authors present the results of the application of the developed decision support system that uses the LP model, which led to tangible savings for the company.

2.3.3 Short-term planning literature

We conclude the review by presenting the studies on short-term planning problems, specifically production scheduling. Based on a case study at an agrochemical plant in the Middle East, Dessouky et al. (1999) develop a mixed integer nonlinear programming model (MINLP) to solve a short-term production planning and scheduling problem. The plant needs to schedule the production of multiple products produced in a set of resources. Setup times are considered, but assumed to be independent of the production sequence, and all batches are assumed to have an identical processing time. The goal of the model is to simultaneously allocate customers' orders to production batches and determine the production schedule that minimizes earliness and tardiness costs. The authors reformulate the problem as a MILP and additionally propose a heuristic to solve the problem. Finally, they compare the performance of the two solution approaches in a numerical study.

McGraw and Dessouky (2001) extend the model developed by Dessouky et al. (1999) to consider sequence-dependent setup costs. The authors also present different heuristics for different variants of the problem, and subsequently asses their performance.

Batching, scheduling and other similar operational production planning problems have been the subject of many studies with applications in the general process industry; although we do not review this literature in further detail, we note that a significant number of these studies could be relevant for short-term production planning in the agrochemical industry.

To summarize, we see that many of the unique characteristics and challenges of the agrochemical industry mentioned in the previous sections are reflected in the literature on all SC planning levels. First, seasonality of demand is considered in nearly all long-term and mid-term planning problems. In mid-term planning problems, seasonality makes it necessary to consider a planning horizon which lasts at least a complete seasonal cycle, because capacity in off-season periods are required to build up stock for the peak season. In long-term planning problems, demand seasonality must be considered in order to obtain correct and accurate estimates of the effects of network design decisions, such as the capacity of the facilities in the SC network. Second, the complexity of the production process of agrochemicals, especially the need for costly and time-consuming changeover operations, is also considered in all planning levels. Indeed, different authors stress the importance of using detailed production planning models to determine the feasibility of, and provide feedback to, more aggregate higher-level models, in order to obtain realistic and accurate results from the latter. Third, the importance of R&D in the industry has led to the development of detailed models to aid companies in selecting and scheduling research and testing activities. Finally, we note that, in most studies, despite its importance, uncertainty of demand is either disregarded or considered only through sensitivity analyses. As noted by Bassett (2018), the uncertainty of the timing and length of growing seasons in different regions can potentially have a significant effect on the economic performance of an agrochemical company. The stochasticity of the selling season and its impact on the optimal production and inventory decisions of the firm is precisely what we focus on in this thesis. In Chapter 5 we contribute to the literature on SC planning in the agrochemical industry by developing a mid-term inventory planning model that explicitly considers demand timing and quantity uncertainty, and applying it in a real-world inventory analysis case study.

3 Inventory decisions for a stochastic season: a numerical study

3.1 Introduction

In this chapter, we start to analyze the problem introduced in the previous chapters which is central to this thesis – the inventory management problem faced by a firm that manufactures products with a seasonal and uncertain demand in a common make-to-stock production environment. The starting point of our analysis in this chapter is the study by Schlapp and Fleischmann (2020). The authors developed a theoretical model to examine the inventory problem of a firm selling products over a limited selling season, facing uncertainty about the timing of the season, i.e. its start and end, the total demand over the season, and the temporal distribution of this demand over the season. They analytically derived the firm's optimal inventory policy, which consists of an inventory quantity and an inventory timing decision, providing a detailed analysis of the interplay between these two decisions. We contribute to the inventory management literature by complementing the analysis by Schlapp and Fleischmann (2020) with a numerical study of the problem to investigate how and to what extent the parameters of the problem influence the optimal inventory decisions and the profits of the firm.

The analyzed inventory problem is of high practical relevance, as many goods have seasonal demands, and in a make-to-stock setting demand is always subject to some degree of uncertainty. Either climate or convention can drive the seasonality of demand (Mitchell, 1927). Climatic seasons are the source of seasonality in the demand for many goods: for example, demand for ice cream peaks in the hot summer months, whereas demand for hot chocolate peaks in cold winter months. Conventions, which can have many origins, e.g. religious holidays, fashions and business practices, can also be the driver of seasonal variations in demand for many products: for example, the demand for fireworks is highest in the weeks preceding New Year's Eve celebrations, and the demand for confectionery peaks before festivities such as Halloween and Easter.

When demand for its products is seasonal, the company faces two types of demand uncertainty: uncertainty in demand quantity and in demand timing. *Quantity uncertainty* refers to the stochasticity of the total demand for a product during its selling season, whereas *timing uncertainty* refers to the stochasticity of the properties of the actual selling season such as its start and duration. Clearly, the latter is more severe in case the seasonality of demand is due to climate, because weather conditions are difficult to forecast with accuracy. Given these two types of uncertainty, the company must make two related inventory decisions (Schlapp and Fleischmann, 2020): an *inventory quantity* and an *inventory timing* decision, i.e. the quantity to stock for the selling season, and when to make the stock available for sale to the customers, respectively. As shown by

Schlapp and Fleischmann (2020), both decisions affect the company's expected profits, and must be made in conjunction with each other. The larger the inventory quantity, the larger the potential sales, but also the larger the potential leftovers, giving rise to the well-known newsvendor trade-off between underage and overage costs. The timing of this quantity being made available for sale also affects underage and overage costs. Indeed, the later the inventory timing is, the higher the risk of losing early demand in the season and of ending the season with more leftovers; however, importantly, a later timing also decreases the inventory holding costs which are incurred until the products are either sold or salvaged. This, in turn, increases the profit margin and decreases the overage cost per unit of the product, making the inventory timing and quantity decisions dependent on each other.

The agrochemical industry, presented in Chapter 2, is a perfect example of a setting in which demand seasonality is driven by climate, and both the inventory quantity and timing decisions significantly impact the companies' profitability.

Due to its practical relevance, the inventory problem for products with seasonal and uncertain demand has attracted significant attention in the operations research literature. However, although demand quantity uncertainty has been extensively studied, timing uncertainty has received little to no attention despite being of equal importance. The newsvendor model and its extensions studied in the literature mainly provide managerial guidelines for the inventory quantity decision, whereas they neglect the inventory timing decision. Indeed, usually the start of the season is assumed to be known with certainty, and the length of the season and its effects on inventory costs are disregarded. To the best of our knowledge, the only study which considers both decisions and uncertainties simultaneously is Schlapp and Fleischmann (2020). In this chapter, we build on those authors' analysis and conduct a numerical study with the following objectives: First, we investigate how the parameters of the problem affect the two key inventory decisions. Second, we assess the effects on the firm of neglecting, or incorrectly considering, the timing uncertainty when setting its inventory policy.

The remainder of this chapter is structured as follows. In Section 3.2, we conduct a review of the literature most closely related to the inventory problem. In Section 3.3, we present the details of the analyzed problem setting and briefly summarize the properties of the optimal inventory policy of the firm following the results of Schlapp and Fleischmann (2020). In Section 3.4, we present the setup of the numerical study that is performed to answer our research questions, and in Section 3.5 we show and interpret the results. Finally, in Section 3.6, we discuss the main managerial insights obtained and identify further research opportunities.

3.2 Literature review

With the study presented in this chapter, we contribute to the literature on inventory management for products with seasonal and uncertain demand, which is a very rich field of operations research. The central model in this stream of literature is the classic newsvendor model (Arrow et al., 1955). This model determines the optimal, i.e. expected-profit-maximizing, inventory quantity for a single product with a single selling season and uncertain demand. The classic newsvendor model makes many simplifying assumptions, but a large body of literature focuses on extending the basic model. However, one important assumption has received close to no attention in the literature – that of of an instantaneous season (i.e. a season with deterministic length of zero) with a start which is known with certainty. As a result, the model ignores the inventory timing decision and its role in managing the trade-offs that arise under realistic circumstances in which the properties of the season are stochastic. The inventory timing decision and the stochasticity of the season are the focus of the current chapter.

The stream of literature originating from the classic newsvendor model that is most closely connected to the topic of our research is the one studying the benefits of delaying the production/ordering decision in the newsvendor setting. Given that most other assumptions of the newsvendor model are not modified in this literature stream, the main reason for delaying production is to obtain progressively better demand forecasts as the season's start approaches and thereby decrease the quantity uncertainty. However, later production is typically assumed to lead to higher production costs, thus creating a clear trade-off to be managed by the timing of production. These studies can be broadly classified according to the number of lots that can be produced. Under the assumption of a single production lot, whose timing can be dynamically chosen by the firm, Choi et al. (2004) study the potential benefits of delaying production when forecasts improve over time, and Wang and Tomlin (2009) study the same problem under production lead time uncertainty, with later production having a higher degree of lead time uncertainty.

In many other studies, the number of lots is limited to two and, in most cases, the size of the second lot is also limited by capacity constraints (e.g. Zheng et al., 2015). Finally, certain studies consider the possibility of producing more than two lots, e.g. Wang et al. (2012). The main difference between our study and this stream of literature is that, although in the latter the inventory timing decision is considered, the start and duration of the season are assumed to be known. It is considered beneficial to delay production only because improved demand forecasts may be obtained, whereas the role of the inventory timing decision in limiting inventory costs is neglected. Late production is assumed to be more expensive, which could be interpreted as a method to take into account the cost of adding the necessary capacity to produce the desired quantity at a later point in time. However, as recognized by Lau and Lau (1997) in this very stream of research, this assumption might not always be valid, as delaying production can also lead to inventory cost savings. This effect of inventory timing on inventory costs is the central focus of the current chapter. Moreover, in contrast to many studies in this stream of literature, in the problem setting we analyze, the inventory timing decision is completely endogenous.

Another extension of the newsvendor model which is closely related to our research is that of Ravindran (1972). The author presents an inventory model for seasonal products with uncertain and "contagious" demand, i.e. products for which demand early in the season influences demand later in the season, for example due to word of mouth. In this study, importantly, the length of the selling season is modeled as a decision variable. The start of the season is assumed to be known, and its duration positive but known. The decision maker, therefore, decides not only on the production quantity, but also on the length the selling season, i.e. how long to make the product available for sale. A shorter selling season entails a higher risk of losing late demand, whereas a longer selling season might result in high inventory costs and leftovers to be disposed of. In addition, the author develops an algorithm for finding the optimal selling season's length and production quantity. Deciding the length of the selling season can be seen as an inventory timing decision, which is used to manage the trade-offs arising from the fact that, unlike the classic newsvendor model, the length of the season is positive, which makes it necessary to consider the costs of carrying inventory during the selling season. However, there are two important distinctions between this study and ours. Instead of the end of the selling season, in the problem we analyze in this chapter the inventory timing-decision variable is the time at which the product is first made available for sale. Moreover, in our problem setting, the start and duration of the selling season are assumed to be uncertain.

The studies presented thus far are related to our research because they consider the inventory timing decision apart from the classic inventory quantity decision. However, they do not consider the uncertainty in the properties of the season, such as its start and length. Two streams of literature which go in the latter direction are the one on inventory management for products subject to random obsolescence and the one on inventory management for perishable products with a random lifetime. The former stream was initiated by Hadley (1962) and the latter by Nahmias (1977). When there is a certain probability that the product will become obsolete in the future, which will lead to an end of customers' demand (the selling season), the firm producing the product must decide at any given point in time how much inventory to keep in order to balance inventory, shortage and obsolescence costs. A larger inventory will reduce the risk of shortages until obsolescence occurs, however, it will increase the inventory holding costs in the interim, and the risk of incurring obsolescence costs for leftovers once the product becomes obsolete. Similarly, when a perishable product's limited lifetime is random, the producer must decide at any given point in time how much inventory to keep to balance inventory, shortage, deterioration and discarding costs. Both literature streams essentially study a newsvendor-like setting in which the length of the season is uncertain and the start of the season is known: the start of the season is the time at which the production/ordering decision is made, and the length of the season is the time until the product becomes obsolete or reaches the end of its lifetime. However, as opposed to the problem setting analyzed in this study, the inventory timing decision is exogenous.

To the best of our knowledge, Schlapp and Fleischmann (2020) are the only authors to study the role of inventory timing when the properties of the product's season are stochastic. To this end, the authors develop their model setup based on the fundamental elements of the classic newsvendor model. They use a continuous-time framework and characterize the timing-related properties of the selling season with three parameters: the start of the season, the length of the season, and the shape of the customers' demand pattern over the season. In this problem setting, with the additional uncertainty of the total demand over the selling season, they define the inventory policy of the firm using two variables: the inventory quantity and the inventory timing decisions. After modeling the firm's profit-maximization problem, the authors define the necessary optimality conditions that the inventory quantity and timing decisions must satisfy. The former is found to resemble the typical critical fractile (CF) optimality condition of the classic newsvendor model, and the latter reflects the trade-off between inventory costs and lost sales costs. Next, they analyze the relationship between the two decisions and the effect of the inventory timing decision on the firm's service-level performance in detail. Finally, the authors study two special cases, one with instantaneous demand (i.e. the season has a length of zero) and one where the start and length of the season are known with certainty, which help to isolate and demonstrate the effects on the optimal inventory policy of the uncertainty in the season's start and of different possible demand patterns, respectively.

In this chapter of the dissertation, we consider the problem analyzed by Schlapp and Fleischmann (2020), building on their modeling setup and analytical results. We contribute to the literature by conducting an extensive numerical study with the goal of providing clear managerial insights for firms facing this inventory planning problem. Specifically, our contribution is twofold. First, we show and quantify the effect of the different problem's parameters on the optimal inventory policy of the firm. Second, we identify the settings in which a careful inventory timing decision is crucial for the economic performance of the firm, and thus taking the common-practice approach of making the product available as early as possible is especially inappropriate.

3.3 Problem setting and optimal inventory policy

In the description of the problem setting analyzed in this chapter, we follow for the most part the notation, assumptions and model setup used by Schlapp and Fleischmann (2020). The firm under consideration produces a single product with an uncertain seasonal demand, both in terms of quantity and timing. To serve the uncertain demand in the stochastic selling season, the firm needs to decide its production/inventory quantity x, i.e. the quantity to produce and stock ahead of the season, and its inventory timing t, i.e. the earliest time when the inventory quantity x is made available to customers. The firm's inventory policy is thus defined by the pair (x, t). In the model, time is measured continuously. The product's stochastic market environment is characterized by the product's market potential, Q, i.e. the random total cumulative demand in the season, and its non-deterministic selling season. The selling season is described by a collection of three stochastic elements: B, L and $A(\tau|Q, B, L)$, where B is a continuous random variable representing the start of the selling season, L is a continuous random variable representing the length of the selling season, and A is a (deterministic) function that, for each (Q, B, L) tuple, defines the fraction of the market potential of the product that will be realized between time τ and the end of the selling season $B + L$. Q_τ denotes the market potential of the product at time τ, defined as $Q_\tau = A(\tau|Q, B, L)Q$, and by $D_\tau = -Q'_\tau \geq 0$ the stochastic customer demand rate at time τ, where Q'_τ is the derivative of Q_τ with respect to time. The production costs per unit of the product are c, and its selling price is p, and $p > c$. Inventory costs per unit of the product per unit of time are denoted by h. In accordance with the standard newsvendor model, the firm can only produce a single lot of the product. However, this lot may be partially produced within the selling season, because the season's start is uncertain and the season is not instantaneous. Also, the inventory decisions must be made before the quantity and timing uncertainty is realized and cannot be revised at a later point in time. Additionally, production capacity constraints are ignored. Finally, unmet demand is lost and at the end of the season the product has a salvage value of v, with $v < c$.

Before presenting the firm's optimization problem, to better illustrate the model's setup and notation, and to build intuition on the trade-offs faced by the company in its inventory problem, in Figure 3 we depict the development over time of the market potential (solid line) and the inventory position of the firm (dashed line) under two different inventory policies in a given realization of the stochastic selling season. We denote the inventory at time $\tau \geq t$ by $I_{x,t}(\tau) = [x - (Q_t - Q_\tau)]^+$, where $[Z]^+ = \max\{0, Z\}$.

In the setting illustrated in Figure 3a, the firm makes inventory available before the season starts. Additionally, given that the firm chooses an inventory quantity that is lower than the market potential at time t, the availability period, i.e. the interval of time in which customers can buy the product, starts and ends earlier than the selling season. In this setting, the firm carries significant unnecessary *pre-season inventory*,

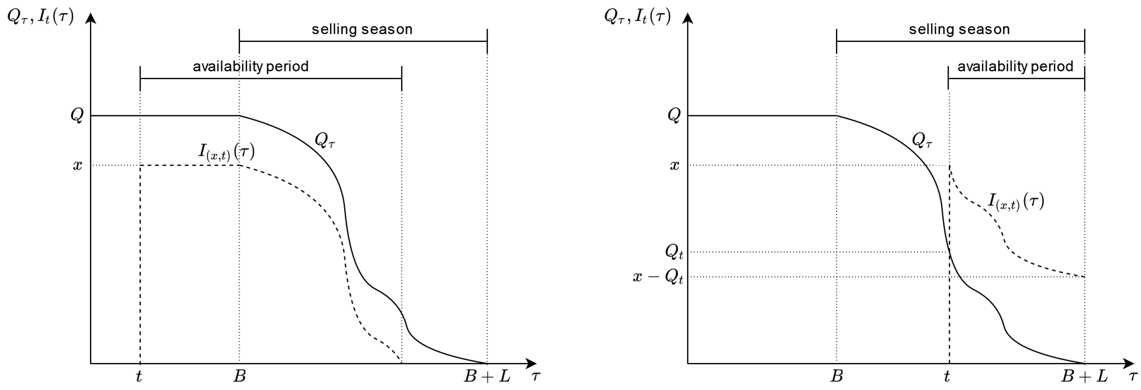

(a) Early inventory timing

(b) Late inventory timing

Figure 3: Illustrative example of two seasons (cf. Schlapp and Fleischmann, 2020)

represented by the area of the rectangle of base $B-t$ and height x. Although the firm is able to satisfy the demand of the earliest customers in the selling season, it loses $Q-x$ demand after the inventory is depleted. The *in-season inventory* is represented by the area under the inventory function to the right of $\tau = B$, and no items must be salvaged at the end of the season. In the example depicted in Figure 3b, the firm instead makes inventory available after the season starts. Because the firm chooses an inventory quantity that is above the market potential at time t, the availability period starts later than and ends at the same time as the selling season. In this setting, the firm carries no pre-season inventory. The late inventory timing leads to a substantial loss of early demand amounting to $Q-Q_t$, and to large left-over inventory at the end of the selling season, equal to $x-Q_t$. These units are needlessly carried in inventory for the whole availability period and are salvaged at the end of the selling season.

The optimization problem faced by the company can be expressed as follows (Schlapp and Fleischmann, 2020):

$$\max_{(x,t)\geq 0} \Pi = \mathrm{E}\Big[(p-c)x - hx(B-t)^+ - h\int_{\max\{B,t\}}^{B+L} I_{(x,t)}(\tau)d\tau - (p-v)(x-Q_t)^+\Big]. \quad (1)$$

The expectation is taken with respect to the joint distribution of B, L and Q. The first and last term within the expectation operator jointly represent the revenue from the sale of the product minus the cost of producing the product and the cost of salvage. The second term represents pre-season inventory holding costs, and the third term the in-season inventory holding costs.

We now characterize the firm's optimal inventory policy following the main results of the analysis presented in Schlapp and Fleischmann (2020). This will serve as a starting point when designing the numerical study reported in this chapter and when interpreting its results.

The optimal inventory quantity decision $x^*(t)$ for any given inventory timing t can be specified as follows. If the expected profit margin of the first unit sold within the availability period is negative, i.e. if $p - c - h\int_t^{b_u} \mathbb{P}(B \geq b)db \leq 0$, then the optimal decision $x^*(t)$ is to set $x^*(t)$ equal to zero; otherwise $x^*(t)$ satisfies the following condition (Schlapp and Fleischmann, 2020):

$$\begin{aligned}
&\mathbb{P}(Q_t \leq x^*(t)) \\
&\quad = \frac{p - c - h[\int_t^{b_u} \mathbb{P}(B \geq b)db + \int_0^{b_u+l_u} \mathbb{P}(Q_t - Q_\tau \leq x^*(t), \max\{t, B\} \leq \tau \leq B + L)d\tau]}{p - v}.
\end{aligned} \tag{2}$$

This condition is closely related to the CF condition of the standard newsvendor model, because the optimal inventory quantity is found by striking the right balance between underage and overage costs. However, two important differences can be readily observed, as noted by Schlapp and Fleischmann (2020). First, the CF relates to the distribution of the market potential (or future demand) at time t, Q_t, and not to that of the total market potential Q. Second, although the CF is still expressed as the ratio of underage costs to the sum of underage and overage costs, these two cost components are different from their counterparts in the newsvendor model due to the presence of inventory costs. Indeed, because the units must be held in inventory until they are sold or salvaged, the profit margin of the units sold decreases, and the revenue of the salvaged units is essentially decreased by the same amount. Consequently, the numerator of (2), representing underage costs, is lower compared to the newsvendor setting, and the denominator, representing the sum of underage and overage costs, is equal to that of the newsvendor's setting due to the decrease in underage costs and increase in overage costs offsetting each other. The net result is that the optimal inventory quantity is lower than the one obtained by applying the newsvendor model to the future demand at time t, Q_t.

Given that the inventory quantity at any time is chosen optimally according to condition (2), and letting $d_\tau(Z) = E[D_\tau 1_{\{Z\}}]$ denote the expected demand rate at time τ in event Z, the optimal inventory timing t^* of the firm can be shown to satisfy the following condition (Schlapp and Fleischmann, 2020):

$$\begin{aligned}
&h\Big(x^*(t^*)\mathbb{P}(t^* \leq B + L) - \int_{t^*}^{b_u+l_u} d_{t^*}(Q_{t^*} - Q_\tau \leq x^*(t^*), \tau \leq B + L, B \leq t^*)d\tau\Big) \\
&\quad = (p - v)d_{t^*}(Q_{t^*} \leq x^*(t^*)).
\end{aligned} \tag{3}$$

The left-hand side of (3) represents the marginal expected (potential) savings in inventory costs obtained by delaying the product's availability period, whereas the right-hand side represents the marginal expected lost revenues caused by losing the immediate demand when inventory timing is delayed. The optimal inventory timing is that which equates these two marginal effects of delaying the availability of the products. More precisely, the first term of the left-hand side measures the marginal savings in inventory costs generated by choosing a later timing, because the optimal inventory quantity must be stored for a shorter time before it is sold. However, choosing a later timing also has an opposite effect on inventory costs, represented by the second term on the left-hand side. The portion of the optimal inventory quantity corresponding to the immediate demand that is lost by delaying the availability period now needs to be stored longer until it is sold. The sum of these two terms, i.e. the entire left-hand side, measures the net effect of a later timing on inventory costs. In contrast, the right-hand side represents the marginal increase in expected lost sales due to the possibility that the aforementioned portion of inventory that is not immediately sold due to a later timing will remain unsold in the remaining part of the selling season.

As can be seen from (2) and (3), the inventory quantity and timing decisions are used to manage two different trade-offs. The inventory quantity manages the trade-off between the classic newsvendor model's underage and overage costs, although both costs are adjusted based on inventory costs. In contrast, inventory timing is used to manage the trade-off between inventory costs and lost revenues. For a more detailed analysis of the problem setting and the optimal inventory policy, readers are referred to Schlapp and Fleischmann (2020).

3.4 Numerical study design

In this section, we present the design of our numerical study in detail. As stated, the goal is twofold: first, we want to analyze how the optimal inventory policy of the firm changes with the parameters of the problem, and second, we want to compare the performance of the optimal inventory policy with that of a naïve classic-newsvendor-like policy which does not optimize the inventory timing decision. We start by defining the parameter settings we consider in our numerical study. Thereafter, we describe the method used to solve the inventory problem instances. Finally, we define the naïve policy used as a benchmark to understand the effects of neglecting the timing uncertainty of demand and the role of inventory timing in managing the trade-offs arising from this uncertainty.

3.4.1 Parameters

In the numerical study, each problem instance solved is defined by a different combination of values of the following parameters:

1. The mean length of the season L;

2. The level of uncertainty in the forecast of the length of the season L;
3. The level of uncertainty in the forecast of the start of the season B;
4. The level of uncertainty in the forecast of the market potential Q;
5. The statistical dependence between the three stochastic parameters defining the season – B, Q and L;
6. The shape of the demand pattern over the season, A;
7. The inventory costs per unit of time and product, h;
8. The critical fractile of the product, CF.

The only two parameters we do not modify across instances of the problem are the means of Q and B, because changing them would effectively only scale up or down the optimal inventory quantity and inventory timing decisions, respectively.

For all numerical parameters, except A and the dependence between B, Q and L, we consider three possible values (or levels, as they are named in ANOVA terminology) that vary across instances. This choice allows us to consider a meaningful range of possible values that the parameter can take while limiting the computational efforts of the study, Moreover, using three levels enables us to verify whether the relationship between the decision variables defining the inventory policy (or the difference in profits between the optimal and naïve policies) and the parameters is monotonic or not. We seek to define parameter levels that lead to problem instances that are practically relevant, and, therefore, offer valid managerial insights.

We now discuss our choice of the levels for each of the parameters listed.

The length of the season

We assume that the stochastic parameter L follows a normal distribution with known mean μ_L and standard deviation σ_L. Due to the symmetric nature of the normal distribution, the normality assumption implies that it is as likely for the selling season to end later than the mean length as it is to end earlier. We believe that this assumption is valid for the purpose of our study and interesting from a theoretical point of view. Naturally, in some practical settings, the symmetric property might not hold, however, the key takeaways derived from our results should be transferable to such settings. We leave the study of other distributional assumptions for future research. We note that, as opposed to Schlapp and Fleischmann (2020), we do not define a lower and upper bound for this parameter in our numerical study.

We assume that the unit of measurement of time is one week, and we consider three possible mean lengths of the season: $\mu_L = \{3, 6, 12\}$. The lowest level corresponds to a short season which lasts less than a month, whereas the highest level corresponds to a season which lasts approximately an entire climatic season. We are of the opinion that this range is representative of selling season lengths. To vary the level of uncertainty in the forecast of the length of the season, we consider three levels for the coefficient of variation of L, $CV_L = \{0.1, 0.3, 0.5\}$, which correspond to three levels of the standard deviation of L for each level of μ_L; these values represent low, medium and high levels of uncertainty, respectively.

We also assume that, for any given realization of the length of the selling season L, customer demand follows the instance-specific demand pattern A such that the total demand over the season is equal to the realization of the market potential Q.

The start of the season

Following the same reasoning described for L, we assume that the stochastic parameter B follows a normal distribution with known mean μ_B and standard deviation σ_B. Also for this parameter, as opposed to Schlapp and Fleischmann (2020), we do not define a lower and upper bound.

As mentioned, the mean start of the season is assumed to be fixed across instances, because its influence on optimal inventory timing is trivial. Because by convention the present is assumed to be $t = 0$, we choose a positive number for this parameter and consider the problem where a firm must decide on its inventory policy for a season which is expected to start μ_B weeks from the present. Specifically, we set $\mu_B = 100$. This choice, combined with a low standard deviation (relative to the mean), avoids possible computational issues with negative values for the potential start of the season. To represent the level of uncertainty in the forecast of the start of the season, we consider three levels for the standard deviation of B: $\sigma_B = \{2, 5, 8\}$. Because time is measured in weeks, we believe that this is a meaningful range of values to consider. The lowest level corresponds to a situation in which the season's start is relatively easy to predict, whereas the highest corresponds to a very unpredictable start of the season. The latter is not uncommon in the agrochemical industry, where the start of a season can shift by as much as two months from one year to the next (see, e.g. Bloemen and Maes, 1992).

The market potential

Similar to B and L, we assume that the stochastic parameter Q follows a normal distribution with known mean μ_Q and standard deviation σ_Q, a common assumption in the newsvendor literature.

We do not vary the mean total demand in the season, because this would only cause a predictable change in the optimal inventory policy, as explained earlier. Therefore, we set $\mu_Q = 100$. As for the other stochastic parameters considered in the numerical study, we vary the accuracy of the forecast of the market potential across instances. Specifically, we consider three levels for the coefficient of variation of Q, $CV_Q = \{0.1, 0.3, 0.5\}$, corresponding to three values of the standard deviation of Q, $\sigma_Q = \{10, 30, 50\}$. We measure Q in general units, because the unit of measurement is inconsequential.

The statistical dependence between the stochastic parameters

In practice, the stochastic parameters of the problem can potentially be correlated with varying degrees of strength, and this correlation could have an important effect on the firm's optimal inventory policy. Therefore, apart from a setting characterized by mutual independence of the parameters (named indep), we consider a single additional setting we believe to be of practical relevance. Specifically, we assume that the pairs (B, L) and (Q, B) are perfectly negatively correlated, and the pair (Q, L) is perfectly positively correlated. In this setting (named pdep), an early start of the season is associated with a longer selling season which accordingly leads to a larger total demand. It is plausible to assume that the end of a climatic season, for example summer, corresponding to the selling season of a product, is uncorrelated with its start, so that the length of the season can depend on its start. In other words, a season that starts early may last longer and still end at the same expected time. If the season lasts longer, it is also plausible that the total quantity demanded would be larger, for example because the value of the product increases if it can be used longer. The correlations in practical settings might not be perfect as assumed, but the results obtained using the assumption of perfect correlation are instructive nonetheless. We denote the statistical dependence between the stochastic parameters by dp, and, as discussed, consider two levels for this parameter, $dp = \{$indep, pdep$\}$. We defer the study of how other potentially practically relevant dependence structures influence the firm's optimal inventory policy to future research.

The shape of the demand pattern

In accordance with studies on inventory problems for items with seasonal demand, we consider four possible seasonal demand patterns (corresponding to four customer demand rates/functional forms of A) in our numerical study: constant, increasing, decreasing and triangular (i.e. first increasing and then decreasing) demand rates. The decreasing and triangular demand patterns are, for example, analyzed by Groebner and Merz (1990) when studying the inventory problem of a retailer of sporting goods facing seasonal demand for its products. The authors state that a decreasing demand pattern may be seen for fishing equipment, while a triangular demand pattern for skis. Similarly, Gupta et al. (2003) develop ordering policies for items with a deterministic seasonal demand which is assumed to be either constant or increasing over the selling season. Although we consider only these four cases, many other plausible demand patterns, such as a trapezoidal demand rate, can be thought of as minor modifications or combinations of the analyzed patterns. We now present the specific model formulation of the analyzed inventory problem for each of these demand patterns; these four sets of formulations - equations (9), (15), (22) and (31) - are special cases of (1).

Constant demand rate

In this case (named const), the assumption is that the demand rate remains constant throughout the entire selling season. For a given scenario s (a realization of the random variables of the problem), we define the cumulative demand function $CD_s(\tau)$, i.e. the cumulative demand that has occurred from the start of the season in scenario s, B_s, to time $\tau > B_s$ as

$$CD_s(\tau) = \frac{Q_s}{L_s}(\tau - B_s). \tag{4}$$

To simplify the notation, in what follows we drop the subscript s. Moreover, to simplify the calculations and for ease of exposition, only for the purpose of calculating the in-season inventory, we translate the season to obtain $B = 0$. Therefore the cumulative demand function is

$$CD(\tau) = \frac{Q}{L}\tau. \tag{5}$$

The market potential at time τ is defined as

$$Q_\tau = CD(L) - CD(\tau) = Q\Big(1 - \frac{\tau}{L}\Big). \tag{6}$$

Figure 4 shows the market potential Q_τ (solid line) and the demand rate D_τ (dashed line) functions for the constant demand rate pattern.

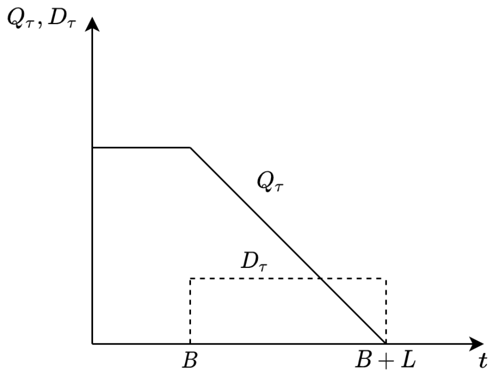

Figure 4: Shape of the demand pattern – constant demand rate

The in-season inventory function for $\tau \in [\max\{t, B\}, B + L]$ can be expressed as

$$I_{(x,t)}(\tau) = (x - (Q_t - Q_\tau))^+ = \Big(x - Q\Big(\frac{\tau - t}{L}\Big)\Big)^+. \tag{7}$$

The total in-season inventory, denoted by $SI_{(x,t)}$, can then be derived as follows:

$$SI_{(x,t)} = \int_l^u I_{(x,t)}(\tau)d\tau = \int_l^u \Big(x - Q\Big(\frac{\tau - l}{L}\Big)\Big)d\tau = x(u-l) + \frac{Q}{L}\Big(-\frac{1}{2}u^2 + lu - \frac{1}{2}l^2\Big), \tag{8}$$

where $l = \min\{L, \max\{0, t - B\}\}$ and $u = \min\{L, \frac{L}{Q}x + l\}$. If $t > L$, then $SI_{x,t} = 0$ because $l = u = L$. If $t < L$, then the lower limit of integration must be equal to t if the product is introduced after the start of the season, and equal to B otherwise. The upper limit of integration must be equal to either the time at which the season ends, L, or the time at which the inventory function reaches zero.

Therefore, for a given inventory policy (x, t) we can calculate the profit of the firm in a given scenario as:

$$\Pi(x,t) =(p-c)x - hx(B-t)^{+} - h\Big(x(u-l) + \frac{Q}{L}\Big(-\frac{1}{2}u^2 + lu - \frac{1}{2}l^2\Big)\Big) \\ - (p-v)\Big(Q\Big(1-\frac{l}{L}\Big)\Big)^{+}. \tag{9}$$

<u>Increasing demand rate</u>

In this case (named incr), we assume that the demand rate is zero at the start of the selling season and increases constantly over the duration of the season. For a given scenario s, the cumulative demand function $CD_s(\tau)$ is assumed to be a parabola facing upward defined as

$$CD_s(\tau) = \frac{Q_s}{L_s^2}(\tau - B_s)^2. \tag{10}$$

Again, in what follows we drop the subscript s and translate the season so that $B = 0$. Therefore, the cumulative demand function becomes

$$CD(\tau) = \frac{Q}{L^2}\tau^2. \tag{11}$$

The market potential at time τ is defined as

$$Q_\tau = CD(L) - CD(\tau) = Q\Big(1 - \frac{\tau^2}{L^2}\Big). \tag{12}$$

Figure 5 shows the market potential Q_τ (solid curve) and the demand rate D_τ (dashed line) functions for the increasing demand rate pattern.

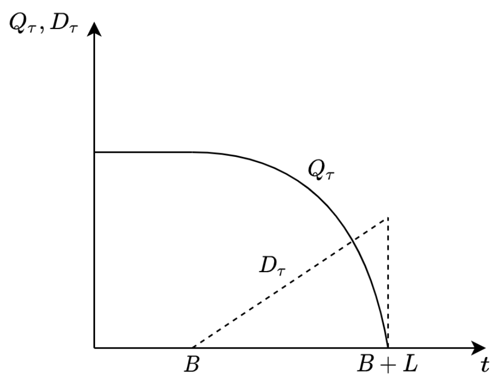

Figure 5: Shape of the demand pattern – increasing demand rate

The in-season inventory function for $\tau \in [\max\{t, B\}, B + L]$ can be expressed as

$$I_{(x,t)}(\tau) = (x - (Q_t - Q_\tau))^+ = \Big(x - Q\Big(\frac{\tau^2 - t^2}{L^2}\Big)\Big)^+. \tag{13}$$

The total in-season inventory, $SI_{(x,t)}$, can be derived as follows:

$$SI_{(x,t)} = \int_l^u \Big(x - Q\Big(\frac{\tau^2 - l^2}{L^2}\Big)\Big)d\tau = x(u - l) + \frac{Q}{L^2}\Big(-\frac{1}{3}u^3 + l^2u - \frac{2}{3}l^3\Big), \tag{14}$$

where $l = \min\{L, \max\{0, t - B\}\}$ and $u = \min\{L, \sqrt{\frac{L^2}{Q}x + l^2}\}$. If $t > L$, then $SI_{(x,t)} = 0$ because $l = u = L$. If $t < L$, then the lower limit of integration must equal t if the product is introduced after the start of the season, and equal to B otherwise. The upper limit of integration must equal either the time at which the season ends, L, or the time at which the inventory function reaches zero.

Therefore, for a given decision (x, t) we can calculate the profit of the firm in a given scenario as:

$$\begin{aligned}\Pi(x,t) =& (p - c)x - hx(B - t)^+ - h\Big(x(u - l) + \frac{Q}{L^2}\Big(-\frac{1}{3}u^3 + l^2u - \frac{2}{3}l^3\Big)\Big) \\ & - (p - v)\Big(Q\Big(1 - \frac{l^2}{L^2}\Big)\Big)^+. \end{aligned} \tag{15}$$

Decreasing demand rate

The assumption in this case (named decr) is that the demand rate is largest at the start of the season and decreases constantly during the selling season until it reaches zero at the end of the season. For a given scenario s, the cumulative demand function $CD_s(\tau)$ is a parabola facing downward defined as

$$CD_s(\tau) = Q_s - \frac{Q_s}{L_s^2}(\tau - B_s - L_s)^2. \tag{16}$$

Again, in what follows we drop the subscript s and translate the season so that $B = 0$. The cumulative demand function therefore becomes

$$CD(\tau) = Q - \frac{Q}{L^2}(\tau - L)^2. \tag{17}$$

The market potential at time τ is defined as

$$Q_\tau = CD(L) - CD(\tau) = \frac{Q}{L^2}(\tau - L)^2. \tag{18}$$

Figure 6 shows the market potential Q_τ (solid curve) and the demand rate D_τ (dashed line) functions for the decreasing demand rate pattern.

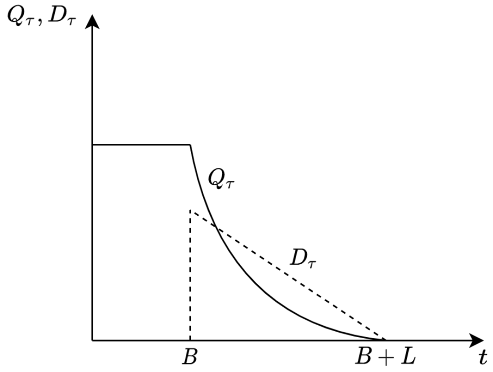

Figure 6: Shape of the demand pattern – decreasing demand rate

The in-season inventory function for $\tau \in [\max\{t, B\}, B+L]$ can be expressed as

$$I_{(x,t)}(\tau) = (x - (Q_t - Q_\tau))^+ = \Big(x - \frac{Q}{L^2}\Big(t^2 - \tau^2 + 2L(\tau - t)\Big)\Big)^+. \tag{19}$$

The total in-season inventory, $SI_{(x,t)}$, can be derived as follows:

$$SI_{(x,t)} = \int_l^u \Big(x - \frac{Q}{L^2}\Big(l^2 - \tau^2 + 2L(\tau - l)\Big)\Big) d\tau = x(u-l) + \frac{Q}{L^2}\Big(-l^2u + \frac{1}{3}u^3 + \frac{2}{3}l^3\Big) - \frac{Q}{L}(u-l)^2, \tag{20}$$

where $l = \min\{L, \max\{0, t - B\}\}$ and u is

$$u = \begin{cases} L - \sqrt{L^2 + l^2 - 2Ll - \frac{L^2}{Q}x}, & \text{if } x < \frac{Q}{L^2}(l-L)^2 \\ L, & \text{otherwise.} \end{cases} \tag{21}$$

If $t > L$, then $SI_{x,t} = 0$ because $l = u = L$. If $t < L$, then the lower limit of integration must equal t if the product is introduced after the start of the season, and equal B otherwise. The upper limit of integration must equal either the time at which the season ends, L, or the time at which the inventory function reaches zero (in this case u is defined with a conditional equation because the inventory function might not reach zero, in which case the square root is undefined).

Therefore, for a given decision (x, t) we can calculate the profit of the firm in a given scenario as:

$$\begin{aligned}\Pi(x,t) =& (p-c)x - hx(B-t)^+ - h\Big(x(u-l) + \frac{Q}{L^2}\Big(-l^2u + \frac{1}{3}u^3 + \frac{2}{3}l^3\Big) - \frac{Q}{L}(u-l)^2\Big) \\ &- (p-v)\Big(\frac{Q}{L^2}(l-L)^2\Big)^+.\end{aligned} \tag{22}$$

Triangular demand rate

In this case (named tri), the assumption is that the demand rate is zero at the start of the season and increases constantly until the middle of the season ($L/2$), and then decreases constantly until the end of the season, where it reaches zero again (we call this case "tri" in short). For a given scenario s, the cumulative demand function $CD_s(\tau)$ is the combination of a parabola facing upward for the first half of the season and a parabola facing downward for the second half of the season (as for the other demand rates, we drop the subscript s and translate the season so that $B = 0$):

$$CD(\tau) = \begin{cases} 2\frac{Q}{L^2}\tau^2, & \text{if } \tau \leq \frac{1}{2}L \\ -2\frac{Q}{L^2}\tau^2 + 4\frac{Q}{L}\tau - Q, & \text{otherwise.} \end{cases} \tag{23}$$

The market potential at time τ is defined as

$$Q_\tau = CD(L) - CD(\tau) = \begin{cases} Q\Big(1 - \frac{2}{L^2}\tau^2\Big), & \text{if } \tau \leq \frac{1}{2}L \\ 2Q\Big(1 - \frac{\tau}{L}\Big)^2, & \text{otherwise.} \end{cases} \tag{24}$$

Figure 7 shows the market potential Q_τ (solid curve) and the demand rate D_τ (dashed line) functions for the triangular demand rate pattern.

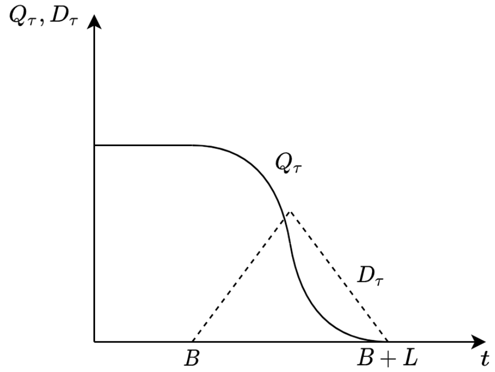

Figure 7: Shape of the demand pattern – triangular demand rate

The in-season inventory function for $\tau \in [\max\{t, B\}, B + L]$ can be expressed as

$$I_{(x,t)}(\tau) = \begin{cases} (x - (Q_t - Q_\tau))^+ = \Big(x - 2\frac{Q}{L^2}(\tau^2 - t^2)\Big)^+, & \text{if } \tau \leq \frac{1}{2}L \\ (x - (Q_t - Q_\tau))^+ = \Big(x - 2\frac{Q}{L^2}(t^2 - \tau^2) + 4\frac{Q}{L}(t - \tau)\Big)^+, & \text{otherwise.} \end{cases} \tag{25}$$

The total in-season inventory for the first part of the season, $SI_{1,(x,t)}$, can be derived as follows:

$$SI_{1,(x,t)} = \int_{l_1}^{u_1} \Big(x - 2\frac{Q}{L^2}(\tau^2 - l_1^2)\Big)d\tau = x(u_1 - l_1) + \frac{Q}{L^2}\Big(-\frac{2}{3}u_1^3 + 2l_1^2 u_1 - \frac{4}{3}l_1^3\Big), \tag{26}$$

where $l_1 = \min\{L/2, \max\{0, t - B\}\}$ and u_1 is $\min\{L/2, \sqrt{\frac{1}{2}\frac{L^2}{Q}x + l_1^2}\}$. If $t > L/2$, then $SI_{1,(x,t)} = 0$ because $l = u = L$. If $t < L/2$, then the lower limit of integration must equal t if the product is introduced after the start of the season, and equal B otherwise. The upper limit of integration must equal either the time at which the first half of the season ends, $L/2$, or the time at which the inventory function reaches zero.

The total in-season inventory for the second part of the season, $SI_{2,(x,t)}$, can be derived as follows:

$$SI_{2,(x,t)} = \int_{l_2}^{u_2} \Big(y - 2\frac{Q}{L^2}(l_2^2 - \tau^2) + 4\frac{Q}{L}(l_2 - \tau)\Big) d\tau = \tag{27}$$

$$= y(u_2 - l_2) + \frac{Q}{L^2}\Big(- 2l_2^2 u_2 + \frac{2}{3}u_2^3 + \frac{4}{3}l_2^3\Big) + \frac{Q}{L}\Big(4u_2 l_2 - 2u_2^2 - 2l_2^2\Big), \tag{28}$$

where y is the leftover inventory from the first half of the season, defined as

$$y = \begin{cases} \max\{0, x - 2\frac{Q}{L^2}\Big(\Big(\frac{L}{2}\Big)^2 - l_1^2\Big)\}, & \text{if } t - B < \frac{L}{2} \\ x, & \text{otherwise,} \end{cases} \tag{29}$$

$l_2 = \min\{L, \max\{L/2, t - B\}\}$, and u_2 is defined as

$$u_2 = \begin{cases} L - \sqrt{(L - l_2)^2 - \frac{1}{2}\frac{L^2}{Q}y}, & \text{if } (L - l_2)^2 > \frac{1}{2}\frac{L^2}{Q}y \iff y < Q_{2,l_2} \\ L, & \text{otherwise.} \end{cases} \tag{30}$$

Therefore, for a given decision (x, t), we can calculate the profit of the firm in a given scenario as:

$$\begin{aligned} \Pi(x,t) =& (p - c)x - hx(B - t)^+ - h\Big(I_{1,(x,t)} + I_{2,(x,t)}\Big) \\ & - (p - v)\Big(y - 2\frac{Q}{L^2}l_2^2 + 4\frac{Q}{L}l_2 - 2Q\Big)^+. \end{aligned} \tag{31}$$

The inventory costs

Both in theory and in practice, it is common to define inventory costs as a carrying charge per unit of product and time as a percentage of the manufacturing costs of the product, c. This is because one of the major components of a product's inventory costs is the opportunity cost of the capital used to manufacture it. Accordingly, we use this parametrization for inventory costs. We denote this carrying charge by $\tilde{h}$ and define different levels of this parameter. This leads to different levels of the original inventory costs parameter of the model, h; specifically, h is defined as $h = \tilde{h}c$. As for the other numerical parameters, we consider three levels of $\tilde{h}$, specifically $\tilde{h} = \{0.1/52, 0.25/52, 0.5/52\}$, representing low, medium and high inventory costs. A range from 10% to 30% per year, and therefore 0.1/52 to 0.3/52 per week, is usually assumed in practice, depending on the application, i.e. the considered product, industry, firm, etc. A high yearly inventory cost level of 50% can apply to products that require special storing conditions, such as certain food products, agrochemicals or pharmaceuticals. Alternatively, this high level can be interpreted as reflecting the decrease in the price of the product as the season approaches its end, which is a common occurrence, for example, in the fashion industry. We note that the different inventory costs levels can also be interpreted as different levels of the unit of measurement of time. For example, assuming that 0.1/52 is the weekly inventory carrying charge, we can interpret the medium and high levels of $\tilde{h}$ as representing a setting wherein the time units of L and B are defined as 2.5 and 5 weeks, respectively.

The critical fractile

We use the classic newsvendor definition of the critical fractile, i.e. $CF = \frac{p-c}{p-v}$, and we consider three levels for this parameter. In accordance with Schweitzer and Cachon (2000), we define products with high and low profit margins as having a CF of 0.75 and 0.25, respectively, and additionally define a medium profit margin product as having a CF of 0.5. Therefore, the three levels considered for this parameter are $CF = \{0.25, 0.5, 0.75\}$. We note that, for all levels, both v and c are kept constant and p is adjusted to obtain the desired critical fractile. This ensures that the CF parameter only controls the profit margin of the product and does not play a role in the determination of the product's inventory costs, which are then controlled by the parameter $\tilde{h}$ alone.

Table 1 summarizes the parameters that are varied across problem instances and the levels considered. We use a full factorial design for our numerical study, solving a total of 5,832 instances of the problem.

Table 1: Parameters of the test instances – inventory management for a stochastic season

Mean length of the season	$\mu_L = \{3, 6, 12\}$
Coefficient of variation of the length of the season	$CV_L = \{0.1, 0.3, 0.5\}$
Standard deviation of the start of the season	$\sigma_B = \{2, 5, 8\}$
Coefficient of variation of the market potential	$CV_Q = \{0.1, 0.3, 0.5\}$
Statistical dependence between stochastic parameters	$dp = \{\text{indep, pdep}\}$
The shape of the demand pattern	$A = \{\text{const, incr, decr, tri}\}$
Inventory costs	$\tilde{h} = \{0.1/52, 0.25/52, 0.5/52\}$
Critical fractile	$CF = \{0.25, 0.5, 0.75\}$

3.4.2 Solution method

Optimization problem (1) is highly complex due to the expectation operator and the nonlinearity of the profit function for a given realization of the stochastic season; no closed-form solution is available, and we are unaware of a method to solve the problem directly. To reduce the complexity of the problem, we first approximate it by representing the uncertainty in the parameters defining the stochastic season – B, L and Q – using a discrete number of scenarios S, a common procedure used in stochastic programming.

Denoting the scenario counterparts of the uncertain parameters and functions of the problem with a subscript s, we can formulate the scenario approximation of problem (1) as

$$\max_{(x,t)\geq 0} \Pi = \sum_{s=1}^{S} \Big[(p-c)x - h\int_t^{B_s+L_s} I_{(x,t,s)}(\tau)d\tau - (p-v)(x-Q_{t,s})^+\Big]pr_s, \tag{32}$$

where pr_s is the probability of occurrence of scenario s, $I_{(x,t,s)}(\tau) = [x-(Q_{t,s}-Q_{\tau,s})]^+$, $Q_{\tau,s} = A(\tau|Q_s,B_s,L_s)Q_s$, and Q_s is the market potential in scenario s.

However, despite this simplification, the problem remains complex to solve. Therefore, for the purpose of the numerical study conducted in this chapter, we use a grid-search procedure to obtain an estimate of the optimal solution to the scenario approximation of the inventory problem of the firm. The basic structure of the grid-search algorithm is outlined in Algorithm 1.

Algorithm 1: Grid-search algorithm

```
Define N grid-points for t, t_n, by dividing the interval [min_s(B_s), max_s(B_s + L_s)]
  into N − 1 intervals of equal length;
Define M grid-points for x, x_m, by dividing the interval [0, max_s(Q_s)] into
  M − 1 intervals of equal length;
t* = 0;
x* = 0;
Π* = 0;
for t = t_1, ..., t_{N−1} do
    Define m^UB = arg min_m x_m | x_m ≥ max_s Q_{t,s};
    for x = x_2, ..., x_{m^UB} do
        Calculate Π = (1/S) Σ_{s=1}^{S} Π_s(x, t);
        if Π > Π* then
            x* = x;
            t* = t;
            Π* = Π;
        end
    end
end
```

In the grid search, we consider $N-1$ equidistant values/points of the variable t in the interval $[\min_s(B_s), \max_s(B_s+L_s)]$, because introducing the product earlier or later is always suboptimal. We then define M values/points of the variable x in the interval $[0, \max_s(Q_s)]$. In the grid search, however, only for $t = t_1$ all M values are considered, because for $t > t_1$ only a subset of them is relevant: $x = 0$ always results in zero profits, so it is unnecessary to evaluate the objective function value at this point; making available a quantity x that exceeds the maximum demand that can be realized before the end of season across all scenarios is also suboptimal, because this will create unnecessary leftovers at the end of the season for every scenario. For a given (x,t) pair, the profit function value is estimated by taking the average of the profit in all scenarios.

The number of points for t, N, and for x, M, and the number of scenarios S influence both the computational time required to solve the problem instances considered and the quality of the solution. Although the computational time increases linearly in all three parameters, the relationship between the quality and the value of the parameters is not as straightforward. To find the right balance between solution time and quality we tested different combinations of values for the three parameters. For each combination, we solved the scenario approximation of the inventory problem and estimated the optimality gap of the solution, as defined in Shapiro (2003). We observed that increasing the value of M and N above 200 resulted in only insignificant improvements in the optimality gap, therefore, we decided to set $M = N = 200$. Of the three parameters, S is the one whose value has the biggest impact on the quality of the solution, as expected in stochastic programming, with the decrease in the optimality gap in S being significant even for large values of S. We took a pragmatic approach and chose to set $S = 150$, which resulted in good quality solutions while limiting the computational effort to a reasonable level.

To obtain the desired number of scenarios S of the three stochastic parameters defining the season for each instance, we use descriptive sampling (DS), first introduced by Saliby (1990), who illustrated the superiority of the method over the standard simple random sampling (SRS) method when solving the classic newsvendor problem. Using DS, in the dp = indep case, we obtain scenarios using the following procedure. First, a set of S values for each stochastic parameter P (with $P = \{B, L, Q\}$) is independently obtained by applying the formula

$$P_s = F_P^{-1}\Big(\frac{s - 0.5}{S}\Big) \qquad s = 1, \dots, S, \tag{33}$$

where F_P^{-1} is the inverse cumulative distribution function of P. Second, for each P, the sampled values are randomly assigned to specific scenario numbers $s = 1, \dots, S$. At the end of this two-step process, each scenario s represents a realization of the stochastic season, with the probability of occurrence of each scenario set equal to $1/S$.

For the dp = pdep case, the first step of the scenario-generation process is the same as for the dp = indep case. However, in the second step, scenarios of the stochastic season are obtained by matching the values of the stochastic parameters obtained in the first step according to their dependence structure.

3.4.3 A naïve inventory policy

As mentioned, to achieve our second research goal, we must define an appropriate benchmark policy, which we name naïve policy. In real-world settings, a common practice is to make the inventory available "as early as possible" to avoid losing sales. Therefore, in each instance of the problem considered, we choose to define the inventory timing of the naïve policy, t^n, as the earliest scenario of B obtained by the DS scenario-generation method in the solution procedure described in Section 3.4.2; formally:

$$t^n = \min_s(B_s). \tag{34}$$

Moreover, given the early timing, we calculate the inventory quantity of the naïve policy, x^n, by applying the classic newsvendor model on the distribution of the (total) market potential Q:

$$x^n = F_Q^{-1}(CF). \tag{35}$$

This naïve inventory policy neglects demand timing uncertainty and the effects of the inventory timing decision on inventory holding costs and, thus, profits.

3.5 Numerical study results

After solving all problem instances with the methods described in Section 3.4, thus obtaining the optimal and naïve inventory policies for each parameter combination, we use ANOVA techniques to analyze and present the results. We divide this section into two parts: in the first part, we describe in detail the methodology used to analyze and visualize the results in order to achieve the research goals set out at the beginning of the chapter; in the second part, we present the results obtained.

3.5.1 Assessment procedure

Our first goal is to determine how the values of the parameters of the inventory problem influence the optimal inventory policy, which is defined by two decisions, the inventory timing and inventory quantity decisions. To this end, for each parameter-decision combination, we first use one-way ANOVA. In this analysis, the dependent (or response) variable is the inventory decision, whereas the independent variable (or factor) is the parameter. ANOVA involves calculating the mean of the response variable for each factor's level over all instances of the problem, and determining the statistical significance of the differences between those means. Formally, ANOVA is used to test the null hypothesis that the mean of the response variable for all levels of the factor is equal against the alternative hypothesis that the mean is different for at least a pair of levels. Defining a significance level of α, an F-test is conducted, and the null hypothesis is rejected if the F statistic obtained from the data is below the threshold value for F corresponding to α. We set $\alpha = 0.05$, which is the significance value commonly used in practice. If the null hypothesis is not rejected, it is concluded that the factor does not have a statistically significant influence on the response variable. In case the null hypothesis is rejected, ANOVA does not provide any information on precisely which means differ from each other. Therefore, to determine which levels have a different mean response, a Tukey's honestly significant difference (HSD) test can be performed, which tests the significance of the difference in the mean of the response variable for all pairs of levels. When informative, we also discuss the results of Tukey's HSD test. In addition, for this test, we use the same significance level as for the ANOVA, i.e. $\alpha = 0.05$.

In our numerical study, multiple factors influence the dependent variables, therefore, a one-way ANOVA is appropriate and informative only if no interaction effects with other factors exist, i.e. if the magnitude and/or direction of the effect of changing the level of one parameter does not depend on the level of other parameters. To assess this, for each parameter-decision combination, we additionally conduct a two-way ANOVA with an interaction effect for each of the other parameters varied in the numerical study. This analysis calculates the mean effect of moving from one level to another of the parameter of interest ("main parameter") for each level of the second ("control") parameter, and then tests the null hypothesis that all these mean effects are equal. To this end, similarly to one-way ANOVA, an F-test is performed (we choose the same significance level as in the one-way ANOVA). We discuss the results of the two-way ANOVAs only if the interaction effect is statistically significant and the interaction modifies the direction of the effect of the main parameter on the inventory decision under consideration, because in this case the results of the one-way ANOVA may be misleading. Moreover, we note that in principle we could also conduct k-way ANOVAs (where k is at most the total number of parameters), as there could be interaction effects between three or more parameters. However, we do not investigate these higher-order interactions, in accordance with the commonly held assumption in practice that these interactions are negligible when the number of factors is moderately large (Montgomery, 2017). For a more detailed explanation of ANOVA techniques, the reader is referred to Montgomery (2017).

To measure the dependent variables in the ANOVAs, i.e. the inventory decisions, we do not use an absolute scale – that is, the unit of measurement of t and x – but a relative one, as we believe this to be more informative. Specifically, for a given instance, we measure the optimal inventory timing as the value of the cumulative distribution function of the start of the season B at t (which is denoted by $R(t)$ to avoid confusion with t), and the optimal inventory quantity as the value of the cumulative distribution function of the total market potential Q at x (which is denoted by $R(x)$); mathematically:

$$R(t) = F_B(t) \tag{36}$$

$$R(x) = F_Q(x). \tag{37}$$

For each factor considered, we also determine which inventory decision changes the most when the level of the factor is changed (in the range of levels considered). This information can provide practitioners with valuable insight on which decision should be managed more closely depending on the characteristics of the product and of the selling season. To this end, for each parameter-decision combination, we first determine the largest and smallest mean value of $R(t)$ over all the levels, denoted by $R(t)^+$ and $R(t)^-$, respectively. We then calculate the difference between these two values, denoted by $\Delta R(t)$:

$$\Delta R(t) = R(t)^+ - R(t)^-. \tag{38}$$

This value measures the strength of the effect that the parameter has on the inventory timing decision. After repeating the same procedure for the inventory quantity decision, we compare $\Delta R(t)$ and $\Delta R(x)$ and determine which inventory decision the parameter affects the most.

Our second goal is to compare the effectiveness of the optimal and naïve inventory policies. To this end, we start by simulating the expected performance of the two policies in each instance of the numerical study. We measure the difference in expected performance between the two policies in terms of relative profit, denoted by $\Delta^r\Pi$ and defined as

$$\Delta^r\Pi = \frac{\sum_{s=1}^{S'} \Pi_s^* - \sum_{s=1}^{S'} \Pi_s^n}{\sum_{s=1}^{S'} \Pi_s^*}, \tag{39}$$

where Π_s^* and Π_s^n are the profits of the optimal inventory policy and the naïve policy in the simulation scenario s in the considered problem instance, respectively, and S' is the number of scenarios used in the simulation. In the simulation we can use a much larger number of scenarios than that used when determining the optimal inventory policy, because the calculation of the simulated expected profits requires an evaluation of the profit function for each scenario for a single inventory policy, (x^*, t^*) or (x^n, t^n), instead of for all the feasible inventory policies considered in the grid-search procedure, described in Section 3.4.2, therefore the computational effort required for each instance is much lower. Moreover, the larger S' is, the more accurate the estimate of $\Delta^r\Pi$ is. We set $S' = 20,000$ in our analysis. Using this value for S', we conducted a test to check the statistical significance of $\Delta^r\Pi$ and found that it was significantly different from zero in

93.11% of the test instances.

We present the results of the performance comparison by focusing on each of the parameters of the problem in turn and show how changing their values from one level to another influences the difference in expected performance of the two policies. This allows us to provide detailed and practically relevant insights. To do this, we apply the ANOVA techniques and procedure, previously described, used to reach the first research goal of this study. In this case, the response variable of the ANOVAs is the difference in expected performance of the two inventory policies, whereas the factors are still the parameters of the problem.

3.5.2 Results

In the results, we focus on one parameter at a time. For each parameter, we present the effect of changing the level of the parameter first for the two optimal inventory decisions, and then for the difference in expected performance between the optimal and naïve inventory policies. We emphasize that our conclusions and interpretations in this section only apply to the specific combinations of levels of the parameters considered in our numerical study.

Mean length of the season – μ_L

Figure 8 shows, for each inventory decision, the main effects plot of the mean length of the season factor on the optimal decision in fuchsia, and, for reference, the naïve policy in black.

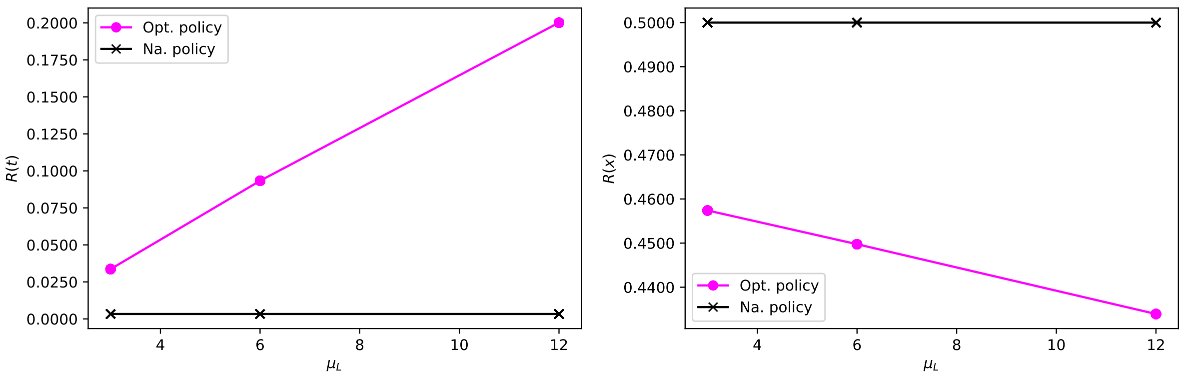

(a) Inventory timing (b) Inventory quantity

Figure 8: Main effects plots – mean length of the season

We note that the scale of the y-axis (the response variable axis) of the plots presented in this section is different for each graph, to ensure that the relationship between the factors and the response variables are clearly portrayed.

Focusing on the inventory timing decision, we see that as the mean length of the season increases, the inventory timing is delayed. Importantly, this relationship is monotonic.

The F-test is highly significant and Tukey's HSD test shows that the difference between all pairs of levels is also significant. The larger the mean length of the season, the higher the in-season inventory costs and, by design, the lower the demand rate in the earlier part of the season. This makes it beneficial to delay timing, because few sales are lost by doing so, however, substantial inventory costs are saved. In contrast, delayed timing in a short season is more expensive, because the firm loses out on a lot more demand in the same time interval, while not achieving significant in-season inventory costs savings.

Considering the inventory quantity decision, as the mean length of the season increases, the optimal inventory quantity decreases, and this relationship is also monotonic. The F-test is significant, but Tukey's HSD test reveals that only the difference in the mean response between $\mu_L = 3$ and $\mu_L = 12$ is significant. A careful examination of the main effects plot reveals that the magnitude of the effect is also negligible. Our interpretation is that the inventory timing decision is very effective in managing the trade-off between early-season lost sales and inventory costs, therefore a significant inventory quantity adjustment is unnecessary. Also, although delaying the inventory timing decreases the future demand, if the length of the season increases simultaneously, the future demand increases by design because the demand rate at each point in time is smaller, thus making it unnecessary to considerably decrease the inventory quantity to account for the reduced future demand.

For this parameter, we obtain $\Delta R(t) = 0.1665 > \Delta R(x) = 0.0234$. The analysis shows that a firm should choose a different inventory policy for the products in its portfolio which have selling seasons of different mean lengths. Specifically, the firm should differentiate between these products by using different inventory timings.

Figure 9 shows the main effects plot for the mean of the length of the season factor on the relative profit difference between the naïve and optimal policies, $\Delta^r\Pi$.

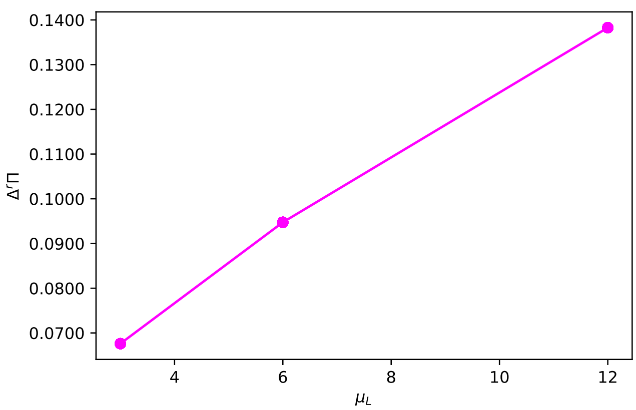

Figure 9: Main effects plot – mean length of the season on relative profit difference

The larger the mean length of the season is, the higher the relative profits difference between the optimal and the naïve policies is. The F-test and Tukey's HSD test for all pairs of levels are statistically significant. As shown in Figure 8a, the optimal inventory timing increases in μ_L, therefore the naïve policy of making the inventory available as soon as possible becomes increasingly inappropriate, leading to a larger relative difference between the two inventory policies. Moreover, the later timing of the optimal policy also leads to a relatively lower future demand and thus a lower optimal quantity compared to the naïve policy, as depicted in Figure 8b.

Coefficient of variation of the length of the season – CV_L

Figure 10 shows, for each inventory decision, the main effects plot of the coefficient of variation of the length of the season factor on the optimal decision in fuchsia, and the naïve policy in black by comparison.

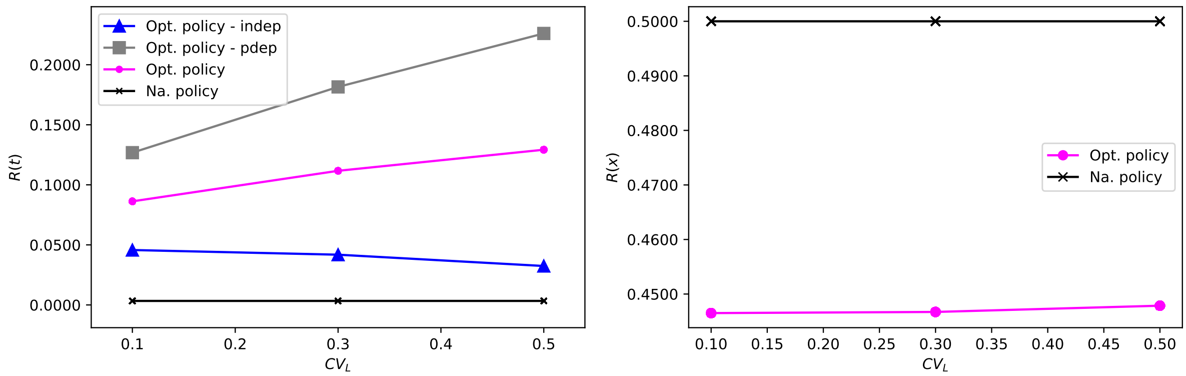

(a) Inventory timing

(b) Inventory quantity

Figure 10: Main effects plots – coefficient of variation of the length of the season

Concerning inventory timing, although the one-way ANOVA and Tukey's HSD test are significant, the conducted two-way ANOVA tests show that there is a significant and important interaction with the factor dp, i.e. the statistical dependence between the stochastic parameters of the problem. Therefore, the results of the two-way ANOVA with the control factor dp are shown in Figure 10a.

The direction of the effect of CV_L on timing depends on the level of dp: specifically, the optimal inventory timing increases in CV_L when dp = pdep and decreases when dp = indep. Our proposed interpretation for this result is as follows. In general, when CV_L increases, the risk of both a shorter and longer season increases. A shorter season is problematic because a late timing is more costly, since by design the demand rate within the season is higher, thus leading to larger lost sales within the same time interval; this makes an early inventory timing more attractive. A longer season is problematic because it causes larger (in-season) inventory costs; this makes late inventory timing more attractive. The net change in the optimal inventory timing caused by an increase in CV_L depends on the ratio of lost sales to inventory costs. In the indep case, the net effect is an earlier timing, due to the lost sales costs being much higher than inventory costs in the problem instances considered. In the pdep case, the consequences of a shorter season are not as problematic as in the indep case, because a short season statistically starts later, has a lower total demand, and has a lower demand rate at its start compared to the indep case. This leads to a smaller decrease in lost sales than in the indep case if the timing is brought forward. An earlier timing, however, leads to a considerable increase in pre-season inventory costs. In contrast, a higher risk of a longer season leads to a substantial increase of in-season inventory costs, thus making later timing more attractive. As a result, the net effect on inventory timing of a higher risk of both a shorter and an earlier season is delayed timing.

Although the main effects graph for the inventory quantity decision shows an increase of the latter in CV_L, this increase is very weak and the F-test shows that it is statistically insignificant. Therefore, we do not further analyze it. Moreover, there is not a significant interaction effect with dp.

Because the main effect of CV_L on the inventory quantity decision is insignificant, and that $\Delta R(t) = 0.0133 > \Delta R(x) = 0.0014$ for the $dp =$ indep case and $\Delta R(t) = 0.0993 > \Delta R(x) = 0.0014$ for the $dp =$ pdep case, we can conclude that a change in CV_L causes a larger change in the inventory timing decision. The analysis shows that a firm should have different inventory timing policies for products with different degrees of uncertainty regarding the length of their seasons. However, the firm should also consider that the optimal timing for each degree of uncertainty depends on the statistical dependence between the start, length and total demand of the season.

Figure 11 shows the main effects plot for CV_L on the absolute profit difference $\Delta\Pi$. We use the absolute difference because the effect of CV_L on the relative profit difference is statistically insignificant, for the given sample size. However, we note that the direction of the effect is the same for both response variables. As for the effect of the parameter on the optimal inventory decisions, the two-way ANOVA tests show that there is a significant interaction effect with the factor dp. Therefore, we add the results of the two-way ANOVA with the control factor dp in Figure 11.

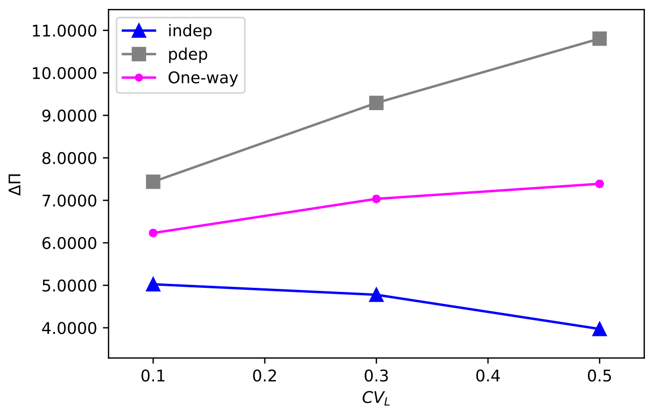

Figure 11: Interaction plot absolute profit difference – coefficient of variation of the length of the season and statistical dependence between the stochastic parameters

The interpretation of the results is based on the difference in timing between the two inventory policies, which can be observed in Figure 10a. The difference in timing, and therefore profits, increases in CV_L for the dp = pdep case, and decreases for the dp = indep case.

Standard deviation of the start of the season – σ_B

Figure 12 shows, for each inventory decision, the main effects plot of the standard deviation of the start of the season factor on the optimal decision in fuchsia, and, for reference, the naïve policy in black.

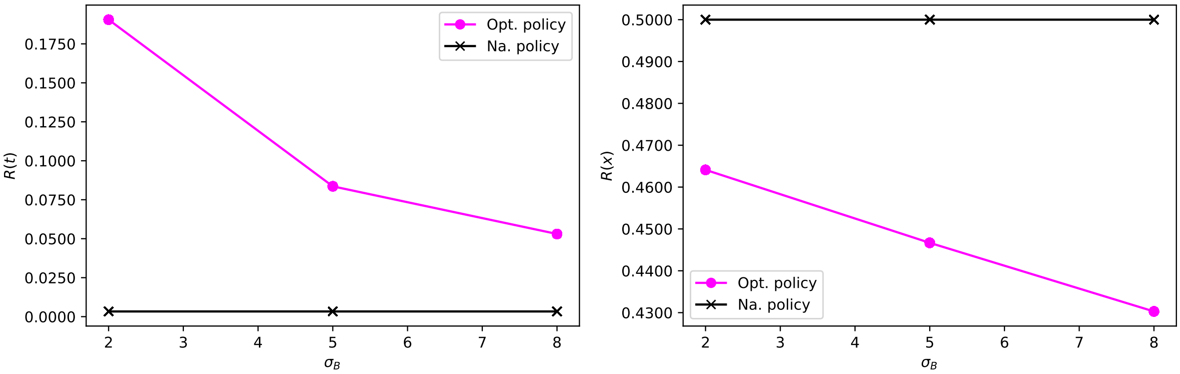

(a) Inventory timing (b) Inventory quantity

Figure 12: Main effects plots – standard deviation of the start of the season

Focusing on inventory timing, the main effects plot shows that the optimal inventory timing decreases in σ_B, and that this relationship is monotonic. The F-test is highly significant and Tukey's HSD test indicates that the difference in the mean optimal timing is significant for all pairs of levels. When σ_B increases, the risk of an earlier and later start of the season are both higher. In the problem instances analyzed in this numerical study, the negative consequences of a late timing, i.e. lost sales, are much more severe than those of an early timing, i.e. inventory costs. As a result, the optimal timing is earlier, to avoid the lost sales from a possible earlier start of the season.

The effect of σ_B on the inventory quantity decision is significant but weak. Tukey's HSD test reveals that only the differences in the mean optimal quantity between the pairs ($\sigma_B = 2$,$\sigma_B = 5$) and ($\sigma_B = 2$,$\sigma_B = 8$) are significant. The optimal inventory quantity decreases in σ_B. Although earlier inventory timing leads to increased future demand, it also leads to increased expected inventory costs per unit, which decrease the profit margin of the product, thus the optimal decision is to slightly decrease the inventory quantity.

For this parameter, we obtain $\Delta R(t) = 0.1375 > \Delta R(x) = 0.0338$. The analysis indicates that a firm should manage the inventory policies of products with different degrees of uncertainty concerning the start of their selling season differently. Specifically, as can be intuitively expected, the inventory policies of such products should be differentiated along the inventory timing dimension.

Figure 13 shows the main effects plot for the standard deviation of the start of the season factor on the relative profit difference $\Delta^r\Pi$.

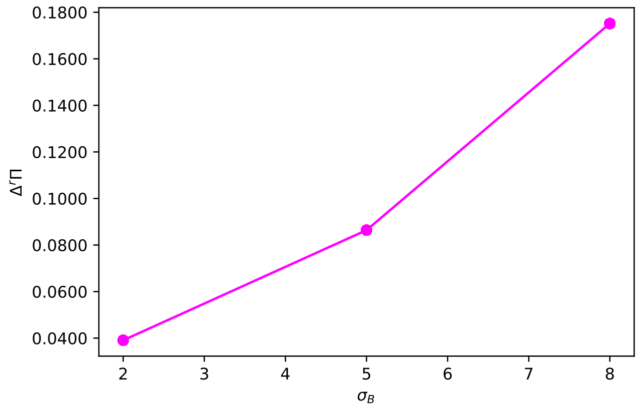

Figure 13: Main effects plot – standard deviation of the start of the season on relative profit difference

The plot shows that the relative profit difference increases monotonically in σ_B. The F-test and Tukey's HSD test are both significant. Because that the optimal timing, measured by $R(t)$, decreases in σ_B, thus decreasing the difference in the timing between the optimal and naïve policies, the direction of the effect of σ_B at first glance appears to be counterintuitive. However, a more careful analysis of the results reveals that, although the difference in $R(t)$ between the optimal and naïve policies decreases in σ_B, the absolute difference in timing t between them increases (this happens because of the larger standard deviation of B), with the optimal policy always being delayed inventory timing. Therefore, the relative profit difference increases in σ_B, due to the larger inventory costs savings achieved by the delayed timing of the optimal policy.

Coefficient of variation of the market potential – CV_Q

Figure 14 shows, for each inventory decision, the main effects plot of the coefficient of variation of the market potential factor on the optimal decision in fuchsia, and the naïve policy in black by comparison.

Considering the timing decision, the optimal inventory timing increases monotonically in CV_Q. Both the F-test and Tukey's HSD test are significant. Our intuition for this result is as follows. For any given timing, the absolute optimal quantity increases if $CF > 0.5$ and decreases if $CF < 0.5$, since the optimality condition for the inventory quantity decision is a newsvendor-like condition adjusted for inventory costs, as shown in eq. (2). If $CF > 0.5$, a larger optimal quantity for any timing should lead to a later inventory timing, because a delay allows to save inventory costs for a larger quantity. This is indeed what the main effects plot shows. In contrast, if $CF < 0.5$, a lower optimal quantity for any timing should lead to earlier timing, because a delay allows to save inventory costs for a lower quantity and, simultaneously, increases expected lost sales. However, this is not what we observe in the results. The two-way ANOVA with

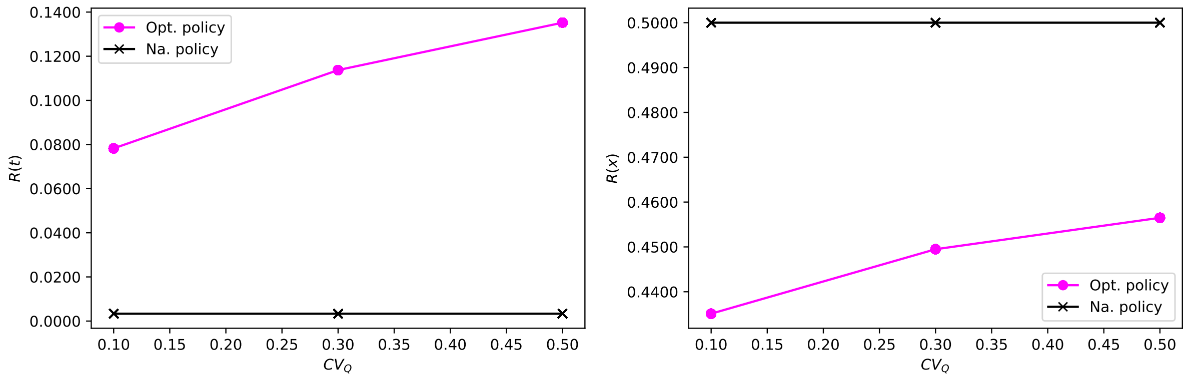

(a) Inventory timing (b) Inventory quantity

Figure 14: Main effects plots – coefficient of variation of the market potential

the interaction effect between CV_Q and CF shows that the direction of the effect of CV_Q on the optimal inventory timing does not depend on the level of CF, i.e. it is optimal to have a later inventory timing even when $CF < 0.5$. We leave this interpretation open for future research to examine.

The main effects plot for the inventory quantity shows an increase in the quantity in the coefficient of variation of the market potential. However, the effect is weak and Tukey's HSD test reveals that only the difference in the mean optimal quantity between the lowest and highest level, the pair $(CV_Q = 2, CV_Q = 5)$, is significant.

We obtain $\Delta R(t) = 0.0570 > \Delta R(x) = 0.0214$. This indicates that the inventory policy of products with a different degree of uncertainty of their total demand over the selling season should be managed differently. Similar to the standard newsvendor model, and thus, the naïve policy, the optimal absolute inventory quantity changes in CV_Q, but the percentile rank of the optimal quantity remains relatively constant. However, the inventory timing differs for different levels of CV_Q, emphasizing the importance of this often neglected component of the inventory policy.

Figure 15 shows the main effects plot for the coefficient of variation of the market potential factor on the relative profit difference $\Delta^r\Pi$.

The results show that the relative profit difference increases monotonically in CV_Q, with both the F-test and Tukey's HSD test being significant. This increase follows directly from the difference in the inventory timing between the optimal and naïve policies, which increases in CV_Q, as shown in Figure 14a. The optimal inventory timing increases, whereas the naïve timing is always the earliest possible one.

Statistical dependence between the stochastic parameters – *dp*

Figure 16 shows, for each inventory decision, the main effects plot of the statistical dependence between the stochastic parameters factor on the optimal decision in fuchsia, and, for reference, the naïve policy in black.

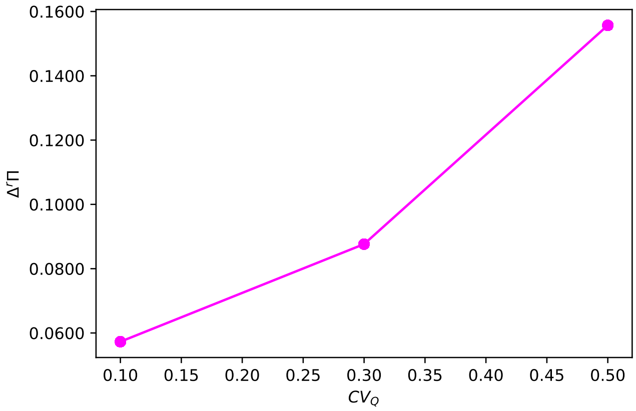

Figure 15: Main effects plot – coefficient of variation of the market potential on relative profit difference

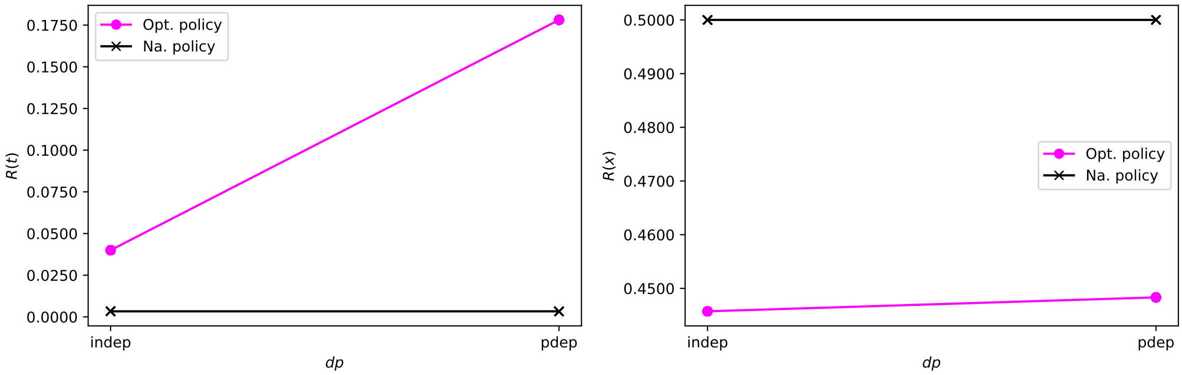

(a) Inventory timing

(b) Inventory quantity

Figure 16: Main effects plots – statistical dependence between the stochastic parameters

The main effects plot for timing shows that the optimal timing in the pdep case is substantially larger than in the indep case. The F-test indicates that the difference between the two means is highly significant. In the pdep case, starting from a given inventory timing, further decreasing the timing reduces the lost sales as in the indep case, but this reduction is not as strong: an earlier timing reduces lost sales costs in the scenarios in which the season starts early; however, in these scenarios, the demand rate at the start of the season is lower due to the longer season, and the future demand is higher due to the extended season, therefore, the earlier timing does not result in many additional sales. Instead, the increase in inventory costs caused by the earlier timing is more substantial, because in these scenarios, as mentioned, the season lasts longer. As a result, in the pdep case, for a given timing, a marginal decrease in inventory timing is less attractive than in the indep case, leading to delayed inventory timing being optimal.

The main effects plot for quantity shows that the inventory quantity is larger in the pdep case, but this difference is small and statistically insignificant. Consequently, we do not further discuss the effect of the statistical dependence between the stochastic parameters defining the season on the optimal inventory quantity decision.

Given that the main effect of dp on the inventory quantity decision is insignificant, and that $\Delta R(t) = 0.1381 > \Delta R(x) = 0.0026$, we can conclude that a change in dp causes a larger change in the inventory timing decision. As a result, a firm producing products with different dependence structures between the stochastic parameters defining the season should differentiate the inventory policies of such products in terms of timing.

Figure 17 shows the main effects plot for the statistical dependence between the stochastic parameters factor on the relative profit difference $\Delta^r\Pi$.

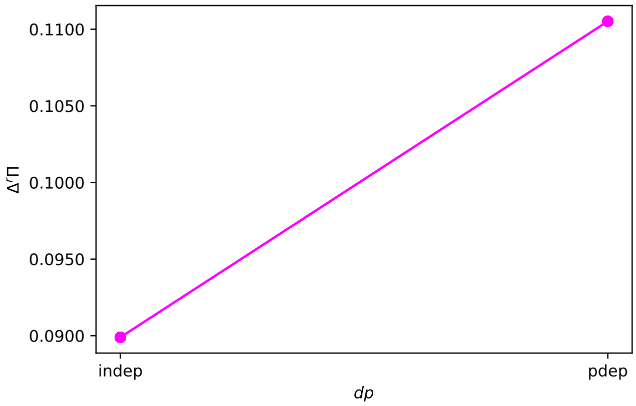

Figure 17: Main effects plot – statistical dependence between the stochastic parameters on relative profit difference

The main effects plot shows that the relative profit difference is higher in the pdep case, with the F-test being significant. Also in this case, the difference in expected relative profits follows directly from the difference in timing between the optimal and naïve policies: indeed, as can be seen in Figure 16a, this inventory timing difference is larger in the pdep case.

The shape of the demand pattern – $\boldsymbol{A}$

Figure 18 shows, for each inventory decision, the main effects plot of the shape of the demand pattern factor on the optimal decision in fuchsia, and the naïve policy in black by comparison.

Concentrating on the inventory timing decision, the F-test indicates that, as the main effects plot suggests, the demand rate has a statistically significant effect on the optimal inventory timing. Tukey's HSD test shows that the difference in optimal timing for all pairs of levels, apart from (const,tri), is significant. The significance of the effect clearly confirms the results of the theoretical analysis presented by Schlapp and Fleischmann (2020), who conclude that the shape of the season plays an important role in determining for how long the inventory timing should be postponed.

For the decreasing, increasing and triangular demand rate cases, there is a clear peak in demand, as illustrated in Figures 5 to 7. In these settings, it is advantageous to make the inventory available at the peak demand time, because this allows the firm to avoid losing

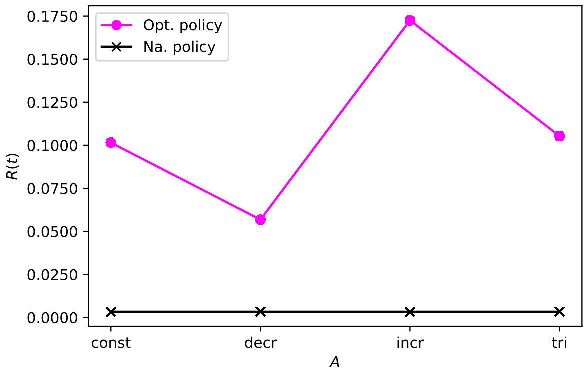

(a) Inventory timing

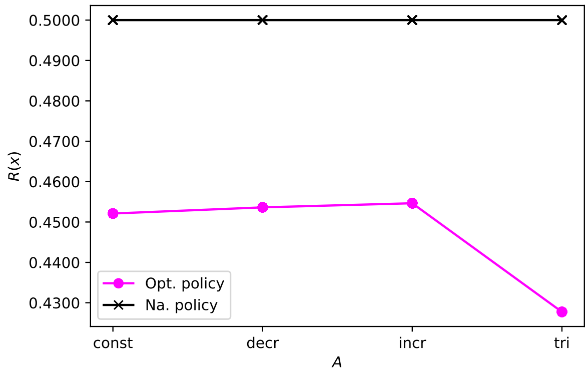

(b) Inventory quantity

Figure 18: Main effects plots – the shape of the demand pattern

a significant amount of potential sales, and to get rid of a large portion of inventory quickly, since the higher demand rate leads to more products leaving inventory quicker. The optimal inventory timing is the earliest in the decreasing demand rate case and the latest in the increasing case. Earlier timing is preferred if there is a decreasing demand rate because the peak in demand happens very early, whereas with an increasing demand rate, a later timing is preferred because the peak happens very late in the season. When the demand rate is triangular, demand peaks in the middle of the season, therefore the optimal inventory timing is between the two extreme cases of decreasing and increasing demand rates. The constant demand rate case is the only one in which there is no peak in demand to take advantage of. Our intuition for the optimal timing being between the decreasing and increasing demand rate ones, similar to the triangular rate, is that, once the optimal portion of demand to be served is determined, it is beneficial to serve this portion of the demand in the later part of the season to save inventory costs. This way the products produced in excess of demand, that thus would remain unsold, are salvaged immediately, rather than remaining in inventory waiting for the season to end.

The main effects plot for quantity shows that the effect of A on the inventory quantity decision is rather weak and that, as confirmed by the F-test and Tukey's HSD test, the only statistically significant differences in quantity are the ones between the triangular demand rate and each of the other demand rates. Specifically, for the triangular case, the optimal inventory quantity is smaller than for all the other levels of A. Our intuition for this result is as follows. It is beneficial to delay inventory timing to avoid inventory costs in the early slow start of the season, and, simultaneously, it is beneficial to have an earlier timing to avoid lost sales and in-season inventory costs in the slow end of the season. This complex trade-off cannot be managed by only adjusting the timing; quantity must also be adjusted. Specifically, quantity is used to avoid in-season inventory costs in the slower end of the season, whereas timing is mainly used to manage the trade-off between pre-season inventory costs and lost sales. Future studies could conduct a more in-depth and systematic analysis into the reasons for this result, as well as the reasons why the optimal quantity for the constant, decreasing and increasing demand rates are approximately equal.

Comparing the effect of A on the two inventory decisions, we obtain $\Delta R(t) = 0.1157 > \Delta R(x) = 0.0269$. Again, we see that when managing products with a different shape of the demand pattern, a firm should mostly use the inventory timing decision to differentiate between the respective optimal inventory policies.

Figure 19 shows the main effects plot for the shape of the demand pattern factor on the absolute profit difference $\Delta\Pi$. We use the absolute profit difference because the difference in the mean of the relative profit difference is not statistically significant for several pairs of levels of A.

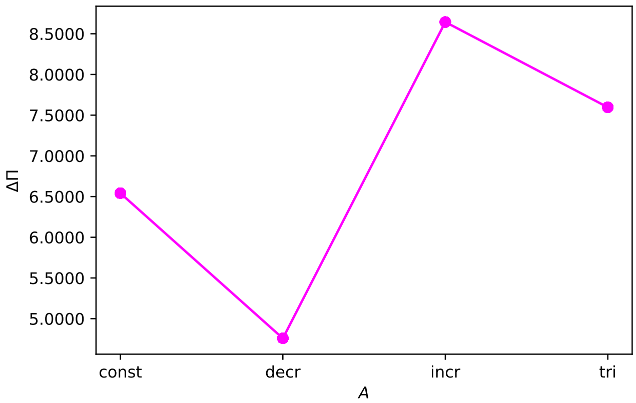

Figure 19: Main effects plot – shape of the demand pattern on absolute profit difference

Both the F-test and Tukey's HSD test are significant. The plot mirrors the one portraying the main effects of A on inventory timing, which provides a clear interpretation of the results. However, it is important to note that, although the difference in inventory timing between the constant and triangular demand rate levels is statistically insignificant, the difference in the absolute profit difference between this pair of levels is significant. The reason is that the optimal quantity for the triangular demand rate case is lower and further away from the suboptimal naïve quantity than the one for the constant demand rate case, thus leading to a larger difference in profits between the two policies.

Inventory costs – $\tilde{h}$

Figure 20 shows, for each inventory decision, the main effects plot of the inventory costs factor on the optimal decision in fuchsia, and, for reference, the naïve policy in black.

The figure shows that as $\tilde{h}$ increases, the optimal inventory timing monotonically increases. The F-test is significant and Tukey's HSD test shows a significant difference between all pairs of levels. The higher $\tilde{h}$ is, the more expensive pre- and in-season inventory is, which makes it increasingly beneficial to accept more expected lost sales to save on inventory costs by delaying timing.

For the quantity decision, as $\tilde{h}$ increases, the optimal inventory quantity decreases mono-

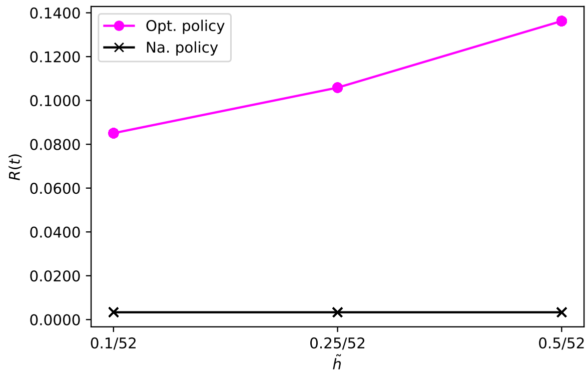

(a) Inventory timing

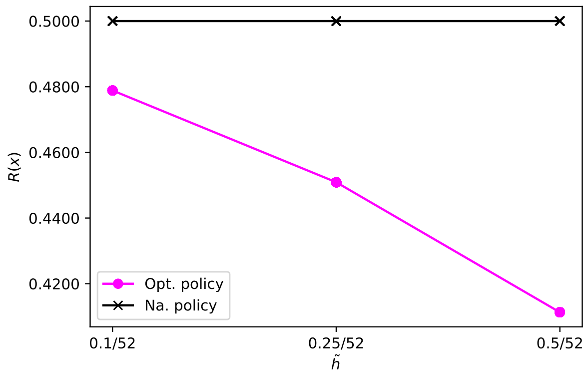

(b) Inventory quantity

Figure 20: Main effects plots – inventory costs

tonically. Also in this case, both the F-test and Tukey's HSD test are significant. When $\tilde{h}$ increases and the timing is delayed, both the future demand and the unit profit margin of the product decrease, thus leading to a smaller optimal inventory quantity. Moreover, Figure 20b confirms the analytical results obtained by Schlapp and Fleischmann (2020) that the optimal inventory quantity is always below the newsvendor one, even with low holding costs, due to the latter model neglecting the complex effect of inventory costs on both underage and overage costs.

The effect of $\tilde{h}$ on the two inventory decisions is similar in strength, with the effect on quantity being slightly stronger: $\Delta R(t) = 0.0512 < \Delta R(x) = 0.0676$. This shows that the firm should have different inventory policies for products with different inventory holding costs and that these optimal policies should differ both in terms of timing and quantity. Moreover, we note that increasing $\tilde{h}$ beyond the range considered in this numerical study, for products with low CFs, the optimal decision in some instances is to not produce/introduce them at all, because all the inventory timing and quantity combinations lead to expected losses.

Figure 21 shows the main effects plot for the inventory costs factor on the relative profit difference $\Delta^r\Pi$.

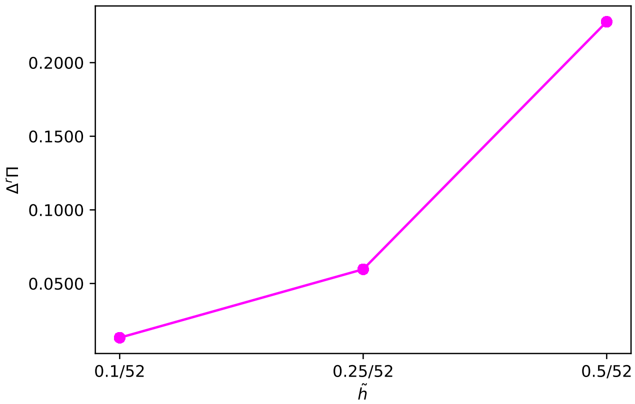

Figure 21: Main effects plot – inventory costs on relative profit difference

The plot shows that the relative profit difference increases monotonically in $\tilde{h}$. The F-test and Tukey's HSD test are both significant. This increase follows directly from the combined difference in the inventory timing and quantity decisions between the optimal and naïve policies. The naïve policy makes the product available as soon as possible to avoid lost sales at all costs. For products with low inventory costs, this strategy works relatively well, as shown in Figure 21. For products with higher holding costs, instead, it is crucial to consider the high inventory costs incurred because of the long time that the products remain on stock before being sold if they are made available too early, which outweigh the expected additional revenues gained if the season starts early.

Critical fractile – CF

Figure 22 shows, for each inventory decision, the main effects plot of the critical fractile factor on the optimal decision in fuchsia, and the naïve policy in black by comparison.

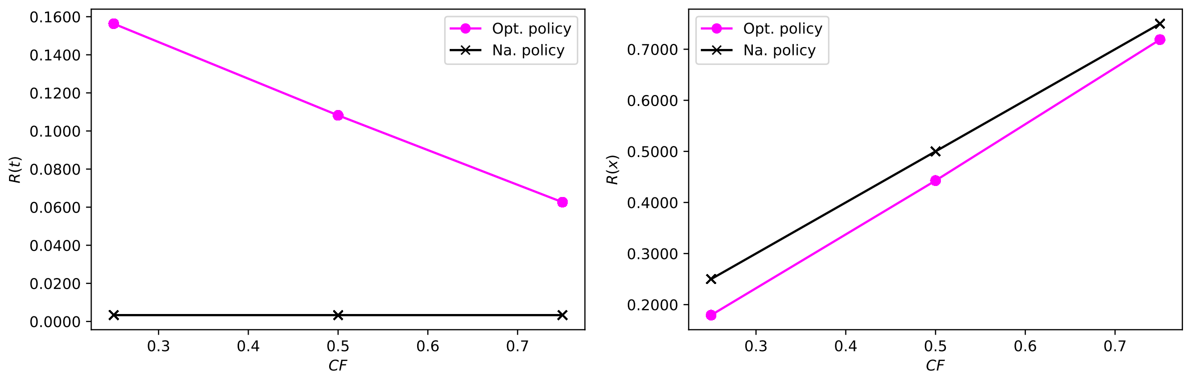

(a) Inventory timing (b) Inventory quantity

Figure 22: Main effects plots – critical fractile

The main effects plot for timing shows that the inventory timing monotonically decreases in CF. The F-test and Tukey's HSD test are significant. The larger CF is, the larger is the ratio of lost sales costs to unit inventory costs is. This makes it beneficial to set an earlier timing to decrease total lost sales, at the expense of increasing expected total inventory costs.

For the quantity decision, as CF increases, the optimal inventory quantity increases monotonically. The F-test and Tukey's HSD test are both significant. With a larger CF and an earlier timing, both the future demand and the unit profit margin of the product increase, thus leading to a larger optimal inventory quantity.

We obtain $\Delta R(t) = 0.0937 < \Delta R(x) = 0.5394$. The analysis shows that products with a different CF should be managed differently, in terms of inventory timing and, more importantly in terms of inventory quantity.

Figure 23 shows the main effects plot for the critical fractile factor on the relative profit difference $\Delta^r\Pi$.

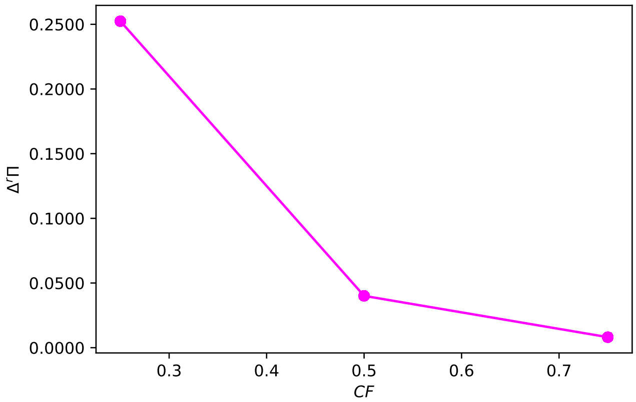

Figure 23: Main effects plot – critical fractile on relative profit difference

The relative profit difference decreases monotonically in CF, and the F-test and Tukey's HSD test are significant. The difference in both inventory timing and inventory quantity between the optimal and naïve policies decreases in the critical fractile, which explains the relative profit differences. Figure 22b again shows that the optimal inventory quantity is always below the newsvendor one, although the difference is small. However, the large gap between the inventory timing of the two policies depicted in Figure 22a shows that the naïve policy, aimed at avoiding lost sales at all costs, is clearly suboptimal when the unit profit margin of the product is low, because the high inventory costs caused by this strategy have a substantial negative impact on the expected profits of the firm.

3.5.3 Synthesis

We summarize the results of the numerical study in Table 2. For each numerical factor, in the *Direction* columns we use plus and minus signs to represent whether changing the value of the factor from its lowest to its highest level respectively increases or decreases the optimal inventory timing and quantity, or the absolute or relative expected profit difference between the optimal and naïve inventory policies. For the non-numerical parameters, in the same column, we report the levels of the parameters in ascending order in terms of the value of the corresponding mean response variables. For the two inventory decisions, we then repeat the values of $\Delta R(t)$ and $\Delta R(x)$, reflecting the strength of the effect that the factor has on the decisions. For the profit difference, we additionally report the indicator $\Delta\Pi^{r+}$, defined, similarly to $R(t)^+$ and $R(x)^+$, as the mean value of $\Delta\Pi^r$ when the factor is at the level which achieves the highest mean value of $\Delta\Pi^r$. This indicator, when compared to the mean of $\Delta\Pi^r$ over all instances of the problem considered in the numerical study, equal to 0.10, facilitates the interpretation of the results of the analysis related to our second research goal.

For the range of values of the parameters of the problem considered in this numerical study, the results show that the effect of changing the level of a parameter on the

Table 2: Numerical study results – summary

Parameter	Inventory timing		Inventory quantity		Profit difference	
	Direction	$\Delta R(t)$	Direction	$\Delta R(x)$	Direction	$\Delta \Pi^{r+}$
μ_L	+	0.1665	-	0.0234	+	0.1383
CV_L	dp=indep: - dp=pdep: +	0.0133 0.0993	+	0.0014*	dp=indep: - dp=pdep: +	0.1018
σ_B	-	0.1375	-	0.0338	+	0.1752
CV_Q	+	0.0570	+	0.0214	+	0.1557
dp	(indep,pdep)	0.1381	(indep,pdep)	0.0026*	(indep,pdep)	0.1105
A	(decr,const, tri,incr)	0.1157	(tri,const, decr,incr)	0.0269	(decr,const, incr,tri)	0.1207
$\tilde{h}$	+	0.0512	-	0.0676	+	0.2278
CF	-	0.0937	+	0.5394	-	0.2524

*Statistically insignificant ($P > 0.05$)

optimal inventory timing of the firm is statistically significant for all the parameters of the problem, whereas the effect on the optimal inventory quantity is insignificant for two parameters, i.e. the coefficient of variation of the length of the season and the statistical dependence between the stochastic parameters defining the season. The median strength of the effect that a parameter has on the inventory timing decision, measured by $\Delta R(\cdot)$, is 0.105, whereas that on the inventory quantity decision is 0.025. This clearly highlights the importance of the often neglected timing component of the inventory policy for products with a stochastic selling season. Additionally, we see that the effect on timing is stronger than that on quantity for all parameters defining the characteristics of the stochastic selling season, i.e. μ_L, CV_L, σ_B, CV_Q, A and dp, whereas the opposite is true for the two cost parameters of the problem, $\tilde{h}$ and CF. This implies that a firm should mainly utilize inventory timing to differentiate the inventory policies of the products in its portfolio when their selling seasons have different characteristics. If these products have different unit profit margins and holding costs, then their inventory policies should also differ significantly in terms of quantity.

In addition, the analysis shows that by correctly considering the effects of inventory costs on profits and the role of inventory timing in managing them, a firm selling products with a stochastic selling season can achieve considerably higher expected profits than if it follows the naïve approach of making inventory available as early as possible. This is especially the case when the start of the products' selling season is highly uncertain, the holding costs are high, and the products have low profit margins.

3.6 Conclusions and outlook

In this chapter, we built on the analysis of Schlapp and Fleischmann (2020), presenting the results of a numerical study on the inventory management problem of a firm that manufactures a product with seasonal and uncertain demand. Using the model setup of Schlapp and Fleischmann (2020), it was possible to differentiate between the two types of demand uncertainty, timing and quantity, confronted by the firm in this setting, and to explicitly model the two key related elements of the firm's inventory policy – the inventory timing and inventory quantity decisions. In our numerical study, we were therefore able to show how and to what extent the parameters of the problem, including the financial ones and the ones describing the timing and quantity uncertainty of demand, influence the two inventory decisions of the firm. Based on our results, we conclude that, although often neglected, inventory timing is a critical component of the inventory policy of the firm, and that it is especially important to differentiate the inventory policies of products with different stochastic selling seasons, as defined by B, L, Q, A and dp, because the strength of the effect of varying the related problem's parameters in the ranges considered, measured by $\Delta R(\cdot)$, is larger on the inventory timing decision than on the inventory quantity decision.

Moreover, by comparing the optimal and naïve inventory policies, we showed that a naïve approach to inventory management for products with stochastic selling seasons is clearly suboptimal. Indeed, importantly, the naïve inventory policy essentially neglects the properties of the product's stochastic season. First, the longer the season is, and hence the less appropriate the assumption of the newsvendor model of an instantaneous season, the more important inventory timing becomes and the larger the relative profit difference between the optimal and naïve policies is. Second, the more uncertain the properties of the season are, i.e. its start, length and total demand, the better the optimal inventory policy performs compared to the naïve one. Third, as already concluded by Schlapp and Fleischmann (2020), the shape of the demand pattern is an additional important characteristic of the season that should be considered when making inventory decisions, because different patterns require different policies, especially in terms of timing. Fourth, the naïve inventory policy disregards inventory costs, and thus the trade-off between lost sales and inventory costs, which is the central trade-off managed by the inventory timing decision. Therefore, when inventory costs are more substantial, considering them becomes crucial for the firm's profitability. This is especially relevant for products with a low CF, because the inventory costs caused by a wrong (early) inventory timing can quickly erode a unit's low profit margin. In settings in which inventory timing is especially important, as just defined, the naïve inventory policy can lead to potential expected losses; these can be up to 250% lower, in absolute values, than the positive expected profits obtained by the optimal policy in the same setting.

The analysis conducted is not without limitations. The quality of the scenario approximation of the analyzed inventory management problem is highly dependent on the number of scenarios used, and the accuracy of the grid-search solution approach is limited. Although the results obtained using this solution approach allow us to draw meaningful and practically relevant conclusions, future studies may be able to improve the precision of the solutions. In addition, due to the computational complexity of solving the problem instances, we can only consider a limited subset of levels for each parameter. Although we believe the levels examined to be practically relevant and the results to be informative for out-of-sample instances, additional parameter settings, for example trapezoidal demand rate and different probability distributions of B, L and Q, can be considered in future research. Moreover, an intuitive interpretation of the effect on one or both optimal inventory decisions is not provided for some of the parameters.

Finally, we note that, although it allows to clearly highlight the effects and importance of demand timing uncertainty and inventory costs on the inventory policy of the firm, the stylized model of Schlapp and Fleischmann (2020) has limited applicability for solving practical inventory problems, which involve many additional complexities not considered by the model. For example, firms usually produce multiple products sharing a limited production capacity, and the production processes are often complicated. In the next chapters of this dissertation, we focus on showing how a firm can efficiently solve its inventory planning problem in a more detailed, realistic setting and consider multiple relevant sources and characteristics of uncertainty.

4 A scenario-generation approach for stochastic lot-sizing problems with correlated demands

4.1 Introduction

In this chapter, we continue our study of the inventory management problem faced by a firm that manufactures products with seasonal and uncertain demands in a common make-to-stock environment. In Chapter 3, using the model developed by Schlapp and Fleischmann (2020), we made some simplifying assumptions which allowed us to focus solely on the effects of seasonality and clearly differentiate between the timing uncertainty and the quantity uncertainty of demand in that setting. For example, we assumed that the product under consideration has a single clear-cut seasonal demand pattern in the planning horizon considered; that at most one lot of the product can be produced; that the production capacity is infinite, with no setup costs and times; and that the firm produces only one product, or, equivalently, that the production capacity of the firm is not shared between the multiple products manufactured. Some of the demand-related assumptions could in principle be dropped by properly modifying the random parameters characterizing the uncertain demand, e.g. by changing the A function used to describe the shape of the season, but to model different and more complex demand patterns it becomes more natural to use a discrete time parametrization of uncertainty. The assumptions related to the production process, instead, require a more complex model than the newsvendor one.

In this chapter, we drop several assumptions made in Chapter 3 and study the same inventory problem in a different production environment and with a more practice-oriented focus. First, we consider a setting in which the firm produces multiple products with a potentially stochastic season that share the limited capacity of a single production resource. Second, we study the problem in an often-encountered production setting in which the production of a product can be started only after a setup operation, with associated setup time and/or setup costs, has been performed on the relevant production resource. Examples of industries where such production processes can be found are the agrochemical, food, beverage, pharmaceutical and semiconductor industries (Copil et al., 2017). Setups make the problem more complex, but at the same time richer by introducing an additional inventory driver, namely economies of scale. However, it should be noted that, although we focus on this specific production environment, the general methodology developed to solve the stochastic inventory problem can be easily applied to other production environments, which can be less or more complex, making the tools presented here relevant for a wide range of applications. Third, we allow products to have different and more complex demand patterns.

Therefore, we model the problem in discrete time and define a demand distribution for each product in each period within the planning horizon. In addition, the demand of different product-period combinations can be dependent. This gives more degrees of freedom to represent demand uncertainty. As a result, we diverge from the previous representation of the seasonal demand uncertainty used in Chapter 3, which uses three random variables (plus a function, which can itself be stochastic) per product, and adopt one that uses one random variable for each product-period combination. This way of representing demand and its uncertainty is very common in the literature and highly relevant in practical applications, because the advanced planning systems employed by many companies use a discrete time axis (Meistering and Stadtler, 2017). To summarize, in this chapter we analyze the problem faced by a manufacturer of products with seasonal and uncertain demand, whose production requires setups. The goal is to determine the optimal production and inventory plan which minimizes costs subject to service-level constraints, as opposed to the goal in Chapter 3, which was maximizing profits. To achieve this objective, the right mix of seasonal, cycle and safety stock must be identified.

To model this problem, we build on the stochastic capacitated lot-sizing problem (SCLSP), the stochastic counterpart of the well-established capacitated lot-sizing problem (CLSP) which has lately received increasing attention in scientific literature. This choice is motivated by the ability of the SCLSP to capture all relevant inventory drivers, i.e. economies of scale, seasonality, and uncertainty, in an integrated fashion. This model also uses a discrete time axis, in accordance with most lot-sizing problems. The SCLSP is commonly modeled as a nonlinear mixed-integer problem, and due to the lack of exact solution methods, approximations are generally developed and solved, possibly with the use of heuristics. The two most notable approximations are the piecewise linear approximation (PLA) and the scenario approximation (SCN). Irrespective of the technique used, it is crucial that the nature of demand uncertainty is preserved as much as possible in the approximated problem, because this determines the validity of the approximation and, therefore, the quality of the solution. Importantly, this means that demand dependencies, that especially characterize settings with seasonal demand, as explained in the introductory chapter of this dissertation, should be correctly reflected in the approximation.

Specifically, when a scenario approximation is used to solve the problem, it is crucial to consider the demand dependencies in the process of generating the scenarios. This is a challenging task for two main reasons. First, in practical applications, *estimating* the correct dependencies is not an easy task. For example, for linear dependencies, an exceptionally large variance-covariance matrix needs to be estimated, possibly from limited or incomplete data. Second, the scenario-generation process must be able to *generate scenarios* reflecting the right dependencies, which calls for techniques that are much more complex to devise and apply than the often used simple random sampling.

In this chapter, we contribute to the stochastic lot-sizing literature by presenting a reformulation of the scenario approximation of the SCLSP, based on a cumulative demand view. This approximation redefines scenarios in a way that significantly simplifies the demand distribution estimation and sampling procedures in the case of demand dependencies, over time and across products.

The remainder of this chapter is structured as follows. In Section 4.2, we present a review of the stochastic lot-sizing literature, followed by a brief review of the literature on scenario-generation techniques. In Section 4.3, we present the nonlinear SCLSP which we use to model the inventory problem studied in this chapter. In Section 4.4, we propose our cumulative scenario approximation of the SCLSP. We then assess the performance of the newly developed approach in Section 4.5. Finally, in Section 4.6, we summarize our contribution and findings.

4.2 Literature review

4.2.1 Stochastic lot sizing

Lot-sizing models are a widely used tool for production planning and inventory management in manufacturing settings in which setup operations are necessary before the production of a product can start. Given that these setup operations entail certain fixed costs which are independent of the quantity produced, these decision models typically focus on striking a balance in the trade-off between setup and inventory holding costs. The fewer the setups performed, the lower the incurred setup costs. However, this implies producing in large lot sizes (or batches) in order to meet the incoming demand, which results in high holding costs. Conversely, performing more setups leads to high setup costs, but low holding costs due to the ability of meeting demand with smaller lot sizes. Therefore, two interlinked decisions must be made to solve these problems: the production timing, i.e. when to produce, and the corresponding lot size.

Lot-sizing problems have been extensively studied, from the simple economic order quantity (EOQ) model to the more complex lot-sizing and scheduling problems with sequence-dependent setup times and/or costs. The EOQ model is a single-item lot-sizing problem with stationary demand, a continuous time axis and an infinite planning horizon (Harris, 1913). All other lot-sizing models can be seen as extensions of the EOQ model. The economic lot-scheduling problem (ELSP) extends the EOQ model by considering multiple products competing for the same scarce capacity (Rogers, 1958). The ELSP keeps the assumptions of a continuous and infinite time scale, which sets it apart from other extensions of the EOQ model, which use a discrete time framework, as well as a stationary demand for the products. Another extension of the EOQ model is the well-known Wagner Within (WW) model, which assumes a discrete and finite planning horizon with dynamic demand, but ignores capacity constraints (Wagner and Whitin, 1958). In the remainder of this chapter, the literature review will focus on extensions of the WW model. These models consider a discrete and finite planning horizon, multiple products with potentially dynamic demand, and capacity constraints. Because the problem setting analyzed in this chapter considers multiple items whose demand can be seasonal, and because of the desired discrete-time representation of demand (and its uncertainty), as explained in the introduction, these models suit our needs very well. According to the way the time horizon is discretized, these discrete-time capacitated lot-sizing models can be broadly classified into small bucket and big bucket models (Eppen and Martin, 1987).

Small bucket problems divide the planning horizon into a large number of short periods, also known as micro-periods. Generally, at most one product can be produced in each of these periods. Problems belonging to this class are the discrete lot-sizing and scheduling problem (DLSP), the continuous setup lot-sizing problem (CSLP) and the proportional lot-sizing and scheduling problem (PLSP). The DLSP makes an "all-or-nothing" assumption, meaning that at most one product can be produced per micro-period and if production occurs, then the entire capacity of that period is utilized (Fleischmann, 1990). The CSLP relaxes this assumption by allowing partial capacity utilization, while still allowing at most one item to be produced per period, meaning that the remaining capacity of the period is unused (Karmarkar and Schrage, 1985). Finally, the PLSP overcomes the drawback of the CSLP that capacity in some micro-periods might be unused by allowing to use this capacity to produce a single additional product in the same micro-period; that is, at most two products can be produced in the same period (Drexl and Haase, 1995).

Big bucket models, instead, divide the planning horizon into a small number of long periods, also called macro-periods. In each of these periods, multiple products can be produced. Big bucket models include the CLSP (Eppen and Martin, 1987) and the capacitated lot-sizing and scheduling problem with sequence dependent setups (CLSD). The CLSP is a basic big bucket model, whereas the CLSD (Haase, 1996) extends it by allowing for sequence dependent setup times and setup carryover.

A model which does not perfectly fit this classification, and therefore is usually referred to as a "hybrid" model, is the general lot-sizing and scheduling problem (GLSP), which divides the time horizon into a finite number of macro-periods, each composed of a fixed number of micro-periods with variable lengths in which at most one product can be produced (Fleischmann and Meyr, 1997). The number of micro-periods restricts the amount of products that can be produced per macro-period.

Most studies on the lot-sizing problem assume a deterministic environment, i.e. all parameters of the problem are assumed to be known with certainty. Buschkühl et al. (2010) provide a review of deterministic capacitated lot-sizing problems, whereas Copil et al. (2017) provide a more recent review on deterministic models that consider lot sizing and scheduling simultaneously. However, over the past few decades, a growing body of research has focused on relaxing the certainty assumption, at least partly, and investigated the lot-sizing problem under uncertainty, thus introducing a new inventory driver into the problem. Apart from exploiting economies of scale, now inventory also serves as a buffer against uncertainty: a larger inventory, although expensive, allows to meet the uncertain demand more reliably. The most recent literature review on stochastic lot sizing by Aloulou et al. (2014) shows that the most commonly considered uncertain parameter is demand, although other uncertainties have also been examined. Recent examples of studies that consider problems with other stochastic parameters are Taş et al. (2018), who examine a CLSP with stochastic setup times, Li and Hu (2017), who consider a CLSD with uncertain demand and workforce productivity, and Hilger et al. (2016), who study a CLSP with uncertain demand and returns of used products. Aloulou et al. (2014) also show that in most cases a stochastic programming (SP) formulation is used to solve the problem (for an overview of SP, see Birge and Louveaux, 2011). This approach assumes that the probability distribution of the uncertain parameters of the problem is known. Recently, however, a limited body of research has focused on robust optimization (RO) modeling techniques (see e.g. Alem et al., 2018). In the remaining part of this section the focus is on stochastic lot-sizing problems with uncertain demand modeled using stochastic programming approaches, because these approaches represent the analyzed inventory management problem the best. This chapter contributes to the literature on this topic.

The modeling approaches used by scholars to solve the stochastic lot-sizing problem found are usually classified using the terminology introduced by Bookbinder and Tan (1988). They classify response strategies to deal with demand uncertainty into three categories: the dynamic, static-dynamic, and static uncertainty strategies. Following the *dynamic uncertainty strategy*, both the production periods, i.e. the periods in which a lot of a product is produced, and the lot sizes may change dynamically every time period and be adjusted in accordance with the realization of the uncertain demand. Under the assumption that capacity is unlimited, Scarf (1959) showed that in a setting with dynamic and uncertain demand the optimal inventory policy is a dynamic order-up-to level (s_t, S_t) policy. Finding the optimal parameters of this policy is a complex task which could potentially be solved using dynamic programming (Tempelmeier, 2013). However, the unlimited capacity assumption prevents the applicability of this model to most real-life production settings, where usually multiple products are produced on the same production resource with limited capacity. When capacity constraints and possibly other complicated factors are taken into account, the problem is typically formulated as a multi-stage SP model, in which each stage, i.e. decision-point, corresponds to a time period in the planning horizon. However, solving this problem is complex, especially because the binary setup variables in each stage are dependent on the realization of the uncertainty up to that stage. Consequently, this problem has been solved only for relatively small instances using scenario trees to represent demand uncertainty.

In the *static-dynamic uncertainty strategy*, the production periods are fixed at the beginning of the planning horizon, whereas the lot sizes can depend on the realization of the uncertain dynamic demand in each period. This problem can be modeled as a multi-stage stochastic program, with each period corresponding to a stage and only the setup decisions being first-stage variables.

The *static uncertainty strategy* assumes that all decisions, i.e. production periods and lot sizes, are fixed at the beginning of the planning horizon. This would be a stochastic program with all production decisions being first-stage decisions. As noted by Tempelmeier (2013), although a dynamic-uncertainty strategy would lead to optimal costs, it would also lead to significant planning nervousness, due to the uncertainty in the timing and size of production lots. To avoid this and keep the problem tractable, many studies in the stochastic lot-sizing literature employ a static uncertainty strategy.

Stochastic lot-sizing problems are commonly formulated as nonlinear mixed integer problems and solution techniques typically use approximations. These approximations rely on linear functions to calculate the statistical expectation of the uncertain elements of the model, e.g. inventory. The accuracy of the approximation of the expectation functions clearly determines the validity of the model and the quality of the solution. The most common approximation technique is to use discrete scenarios to represent the possible realizations of the uncertain parameters. Another popular approach consists of approximating the nonlinear expectation functions using piecewise linear functions.

In what follows, a review of the relevant stochastic lot-sizing literature will be presented. The review, summarized in Table 3, is structured according to the classification scheme of Bookbinder and Tan (1988), starting with the dynamic uncertainty strategy and ending with the static uncertainty strategy.

As already anticipated, most studies using the *dynamic uncertainty strategy* model the stochastic lot-sizing problem as a multi-stage SP and use scenarios to represent demand uncertainty. Brandimarte (2006) is one of the first and most cited studies on stochastic lot sizing, and it is one of the few studies which consider at least one type of demand dependence, namely cross-correlation. The author considers a CLSP under uncertain demand. Uncertainty is modeled using a scenario tree, which is a widely used uncertainty representation method in SP. Unmet demand is assumed to be lost, as opposed to backlogged as commonly done in the stochastic lot-sizing literature, and results in a known cost. This assumption will be discussed in more detail later, when we explain our choice of using backlogs in the inventory problem analyzed in this chapter. The uncertainty strategy employed by Brandimarte (2006) is dynamic, and the problem is formulated using a multi-stage plant-location SP model, with both setup and lot size decisions being scenario-dependent in every stage. Because of the size of the resulting mixed integer program (MIP), the author proposes a heuristic approach to solve it. To generate the scenario tree, different procedures which are able to generate samples from a multivariate normal distribution with correlations are used: pure random sampling (RS), antithetic sampling (AS), Latin hypercube sampling (LHS), and moments matching (MM). In numerical tests, the author solves instances of the problem with stationary, multivariate normal, cross-correlated demand. The study finds that the SP model proposed significantly outperforms a naïve deterministic lot-sizing model based on the expected value of demand, especially when capacity is tight and setup times have a large impact.

Hu and Hu (2018) consider a different production setting with sequence-dependent setup times and costs. They develop a stochastic CLSD under demand uncertainty, where the latter is represented using a scenario tree. The problem is modeled as a multi-stage SP model, in which both production periods and lot sizes can be adjusted in every period of the planning horizon, meaning that the authors essentially utilize a dynamic uncertainty strategy. Unmet demand is assumed to be backordered and is penalized in the objective function. The scenario tree is obtained by first generating a large number of scenarios using MM, and then by decreasing this number using scenario-reduction techniques. These techniques start from a large scenario set and choose which scenarios to keep and which to delete in such a way that the resultant reduced scenario set has a probability distribution as close as possible to the original larger scenario set. In their case study, demands of different product-period combinations are assumed to be independent. The authors show that a multi-stage model results in a 10% cost reduction compared to a two stage model.

Curcio et al. (2018) analyze a GLSP under multi-stage demand uncertainty. The problem is modeled as a multi-stage SP with scenarios, thus following a dynamic uncertainty strategy. Also in this study, backorders are allowed and limited by penalizing them in the objective function, as opposed to by using service-level constraints. The authors focus on developing adaptation and approximate solution strategies, based on both SP and RO methods, that make it possible to obtain good quality solutions to the often computationally intractable multi-stage problem in reasonable time. The solution strategies based on SP use a scenario approach to model the uncertainty of the dynamic demand. The sampling approach used to derive the scenarios is simple random sampling (SRS), which is applicable due to the independence of demand for every product and period. The performance of the different approaches is evaluated in a numerical study, which proves their effectiveness, especially for large instances of the problem.

The final study utilizing a dynamic uncertainty strategy reviewed here is by Chen and Su (2022), who analyze a stochastic CLSD with multiple resources and machine eligibility. Demand is uncertain, and each unit of unmet demand can be backordered at a certain cost. The problem is modeled as a multi-stage SP with scenarios. Three demand patterns are analyzed in their numerical study: dynamic, positive and negative trends. In all cases, the demand for each product-period combination is assumed to be uniformly distributed and mutually independent, and scenarios are obtained through SRS. Their results show that the stochastic model outperforms a deterministic one based on the expected value of demand, especially in the setting in which demand is highly uncertain and has a positive trend.

The majority of the studies on stochastic lot sizing under the *static-dynamic uncertainty strategy* analyzes single-item uncapacitated problem settings (e.g. Rossi et al., 2015). The only paper we have found in the literature which comes closest to utilizing a static-dynamic uncertainty strategy in a multi-item capacitated setting is the one by Hu and Hu (2016), who model a CLSD as a two-stage SP program, where demand uncertainty is represented using scenarios. Introducing overtime, they define all regular time production decisions as first stage variables, and overtime production quantities as second stage variables. In other words, the production periods and the minimum lot sizes are fixed in advance for the entire planning horizon, as in the static uncertainty strategy, whereas the lot sizes in every period can be increased by using overtime. This strategy can be viewed as a special case of a static-dynamic uncertainty strategy. The overtime decisions in each period are assumed to be made under perfect information, i.e. current and future periods' demands are known. The remaining details of the problem are the same as the ones analyzed by Hu and Hu (2018), which is an extension of this problem to a multi-stage setting. The model is tested on a real-life case study. As usual, it is shown that there are benefits to applying the stochastic model as opposed to a deterministic model. Additionally, it is found that the setup costs of the deterministic and stochastic models are very close, which indicates that the production sequence is less sensitive to uncertainty than the production quantities.

The literature on stochastic lot-sizing problems under the *static uncertainty strategy* is much richer, probably because of the advantages of the strategy, which were explained earlier. These advantages are what led us to use this uncertainty strategy in solving the inventory problem studied in this chapter. Tempelmeier and Herpers (2010) first proposed a formulation of the SCLSP with a β service-level constraint because of the practical relevance of the β service-level metric. Unsatisfied demand is assumed to be backordered. The underlying strategy used to deal with uncertainty is the static uncertainty strategy. Given the lack of an exact solution method, the authors proposed a modified ABC heuristic to solve the problem, as opposed to solving it by using a PLA or SCN reformulation. A different heuristic, based on column generation, was later presented by Tempelmeier (2011). The performance of the solution methods was tested using instances with dynamic, normal demand, assuming that there is no auto-correlation and no cross-correlation, as well as with lumpy demand. Results show that, on average, the column generation heuristic outperforms the ABC heuristic.

Helber et al. (2013) presented a nonlinear formulation of a SCLSP with a δ service level, which will be used to limit backlogs also in the study conducted in this chapter. Under the assumption that unfulfilled demand is backlogged, the authors show how this new service level measure is able to reflect both the size of the backorders and the waiting time of customers. The authors use the static uncertainty strategy. In addition, the problem is reformulated using two approximations, PLA and SCN, which aim at linearizing the expected backlog and on hand inventory functions. The resulting MIP problems are then solved using a fix-and-optimize (F&O) heuristic. The heuristic's performance, used for solving both approximations, is assessed using instances with dynamic, normal demand, under the assumption that there are no auto- or cross-correlations. In order to obtain scenarios for the SCN model, the authors utilize both SRS and descriptive sampling (DS). The results show that although the PLA approach outperforms the SCN one, the latter still produces solutions considered robust and, as the authors note, has the advantage of being able to model scenarios where auto- or cross-correlation of demand exists. Moreover, it is shown that the SCN model with DS clearly outperforms the SCN model with SRS.

Tempelmeier and Hilger (2015) subsequently extend the PLA model of Helber et al. (2013) by introducing a β service-level constraint and by providing for setup carry-overs. To solve the problem, they propose a variant of the F&O heuristic, the performance of which is tested and compared against the heuristic presented in Tempelmeier (2011), using the same data set. The newly developed heuristic is shown to outperform the column generation heuristic for smaller problem sizes, whereas it is outperformed by the latter in larger problem sizes. However, the authors point out that, as opposed to the column generation heuristic, the F&O heuristic is able to handle problems which include setup times. Li et al. (2017) analyze the same problem as Tempelmeier and Hilger (2015) but in a setting with multiple resources and overtime production. Moreover, backlogs are controlled by imposing a penalty instead of service-level constraints. The authors reformulate the problem using the PLA approach and develop two new solution approaches based on the F&O heuristic. A numerical study, in which demand is assumed to be stationary, normally distributed and independent over time and across products, is performed to test the newly developed approaches. The authors show that their approaches outperform the F&O heuristic method previously proposed by Sahling et al. (2009). They also compare their solutions to those obtained solving the SCN approximations of the problems, which use SRS to obtain the scenarios and are solved with a branch and cut technique. The results show that their newly developed methods are competitive with the SCN approach for smaller instances, whereas the latter approach is not able to obtain a solution for larger instances. However, this might be because the SCN approach uses a number of scenarios which increases exponentially with the number of periods and that the solution method used is exact instead of being a heuristic.

Alem et al. (2018) investigate a GLSP under demand uncertainty using both SP and RO. In both cases, the static-uncertainty strategy is used to deal with demand uncertainty. Unmet demand is assumed to be backlogged and backlogs are penalized in the objective function. To represent demand uncertainty, the SP approach uses scenarios, which are obtained using SRS given that no dependencies in demand are assumed. The authors compare the two methodologies in computational experiments, confirming their known different properties. Apart from showing that both approaches outperform a deterministic model, the authors also propose guidelines to help decision makers choose the most appropriate uncertainty modeling approach; if possible, given that they each have their own strengths and weaknesses, they should be used together. Moreover, they find that assuming the wrong demand distribution when solving the problem with the SP approach has only a small impact on the performance.

De Smet et al. (2020) study a stochastic CLSD with a β cycle service level under the static uncertainty strategy. The only uncertain parameter considered is demand, which is independently normally distributed for each product and period, and it is assumed that unmet demand is backordered. The authors approximate the problem using the PLA approach and propose a novel procedure to find the most appropriate breakpoints when linearizing the expected backorder and inventory functions. Moreover, they develop a new relax-and-fix with fix-and-optimize heuristic to solve the problem. In a computational study, they prove that their linearization technique leads to cheaper and more conservative plans than previously proposed methods, and that their heuristic outperforms a state-of-the-art solver, both in terms of objective function value and solution time.

We conclude this review of stochastic lot-sizing problems under different uncertainty strategies with three studies that differ from the ones presented so far. These studies use a slightly different approach to model uncertainty and apply their stochastic lot-sizing models under a rolling horizon (RH) planning strategy, which is in itself an alternative method of dealing with uncertainty. The first paper is that of Meistering and Stadtler (2017). The authors analyze a stochastic lot-sizing problem with β cycle service-level constraints. They model the problem starting from a deterministic CLSP, which is then extended to consider demand uncertainty by using exogenous minimum inventory targets at the end of each production cycle (the time between two consecutive setups of the same product), under the assumption that production cycles have integer period lengths. These minimum inventory targets, obtained in a pre-processing phase for all cycles of all possible production schedules, are the safety stocks necessary to achieve the target β cycle service level. Binary variables which determine the start and duration of production cycles of all the products are therefore added to the model, as well as constraints ensuring that whatever schedule is chosen, the inventory at the end of each cycle is larger than the exogenous target. As opposed to the SCLSP, the objective function does not estimate inventory costs based on the expected inventory on-hand, but on the expected net inventory. This method results in an inaccurate estimation of inventory costs, but this inaccuracy is negligible if the target service level is high enough. The focus of the article is not on the specific modeling approach to deal with uncertainty, but on its application in RH planning.

The authors propose the idea of a stabilized-cycle strategy, which can be summarized as follows. At each resolving point in the planning horizon, the estimated future cycle fill-rate is calculated based on currently available information. If this falls below the target, then a cycle that started in the past and finishes in the future is interrupted and a new setup is enforced to avoid missing the target, even if this new setup is made in the frozen horizon of the RH framework. In other words, the problem solved at each re-planning point in time uses a static uncertainty strategy, because decisions are fixed for the entire planning horizon. However, the optimal schedule determined at the beginning of the planning horizon can occasionally be modified at future re-planning points if the developed plan does not meet the intended service-level target. The authors carry out a numerical study comparing their approach to other commonly used strategies to deal with demand uncertainty in an RH planning environment. They analyze both a setting with stationary and seasonal demand, in which demand is period and product independent, finding that their approach reaches a good compromise between inventory costs and the negative deviation of the service-level from its target.

Tavaghof-Gigloo and Minner (2020) study a SCLSP under a RH planning strategy. The authors propose a novel formulation for the SCLSP with cycle service-level constraints which takes into account future re-planning opportunities under a RH strategy. The basic variant of the model uses a static uncertainty strategy. Demand is assumed to be independent over time and across products. Cycle service-level constraints are enforced by ensuring that the inventory level at the beginning of each production cycle is larger than the target inventory level. This target inventory level is a nonlinear function of the mean and the variance of the cumulative demand in the cycle, which are decision variables of the model, because the length of the production cycles is determined endogenously in the model. To linearize the function, a bivariate linearization technique (BLT) is used. Subsequently, the authors present an extension of the model which considers the possible future re-planning opportunities that arise when applying the model under RH planning. This is done by decreasing the target inventory level function at the end of each cycle with the use of a scaling factor calculated in a pre-processing phase, because the re-planning opportunities protect against some demand uncertainty, thus limiting the need for safety stock. As a result, the extended model can be thought of as using a static uncertainty strategy with an exogenous dynamic uncertainty strategy component. The performance of the models is then compared to that of a stochastic dynamic program and a sequential planning approach in a numerical study. The authors find that, as expected, if capacity is limited, the newly developed models outperform the sequential approach, because less inventory is needed to achieve the same service level. If there is sufficient capacity, the extended SCLSP, which considers re-planning opportunities, outperforms the basic SCLSP by avoiding excessive safety stock. Finally, as capacity increases, the performance of the extended SCLSP moves closer to the theoretical lower bound provided by the SDP.

Tavaghof-Gigloo (2019) develops a new formulation of the stochastic GLSP using the same techniques presented in Tavaghof-Gigloo and Minner (2020), that are appropriately modified to fit the different setting that is analyzed. Here backorders are limited by backlog costs as opposed to target inventory levels. This requires the determination of backlogs at the end of each production cycle. In addition, the production cycle lengths are implicit decision variables of the model. The backlog at the end of a production cycle is a nonlinear function of the mean and variance of the cumulative demand within the cycle. The backlog functions are then linearized using a BLT. Moreover, the model can accommodate autocorrelation in the products' demand. Because the model is applied under a RH planning strategy, re-planning opportunities are also considered similar to the 2020 study, namely by decreasing the expected backlog functions at the end of each cycle with the use of a scaling factor calculated in a pre-processing phase. Therefore, the model essentially uses a static uncertainty strategy with an exogenous dynamic uncertainty strategy component. The numerical study conducted by the author on a real-world dataset assumes that demand is independent between products but autocorrelated for each product. The results show that the newly developed approach significantly outperforms other simpler strategies commonly used in practice to deal with demand uncertainty, such as the sequential approach.

Table 3 summarizes the literature presented. The last two columns are only relevant for SCN approaches. In the last column, titles Correlations, the word Yes indicates that the scenario-generation method used to derive the scenarios in the numerical study can accommodate demand dependencies.

Table 3: Stochastic lot-sizing literature

Article	Problem	Uncertainty strategy	Solution method	Scen.-gen. method	Correlations
Brandimarte (2006)	CLSP	Dynamic	SCN	RS, AS, LHS, MM	Yes
Hu and Hu (2018)	CLSD	Dynamic	SCN	MM	Yes
Curcio et al. (2018)	GLSP	Dynamic	RO, SCN	SRS	No
Chen and Su (2022)	CLSP	Dynamic	SCN	SRS	No
Hu and Hu (2016)	CLSD	Static-dynamic	SCN	MM	Yes
Tempelmeier and Herpers (2010)	CLSP	Static	Heuristic		
Tempelmeier (2011)	CLSP	Static	Heuristic		
Helber et al. (2013)	CLSP	Static	PLA, SCN	SRS, DS	No
Tempelmeier and Hilger (2015)	CLSP + carryovers	Static	PLA		
Li et al. (2017)	CLSP + carryovers	Static	PLA		
Alem et al. (2018)	GLSP	Static	RO, SCN	SRS	No
Meistering and Stadtler (2017)	CLSP	static+RH	Exogen. ss targets		
Tavaghof-Gigloo and Minner (2020)	CLSP	static+RH	BLT		
Tavaghof-Gigloo (2019)	GLSP	static+RH	BLT		
De Smet et al. (2020)	CLSD	static	PLA		
This chapter	CLSP	Static	CDS	DS	Yes

CLSP: Capacitated lot-sizing problem, CLSD: capacitated lot-sizing and scheduling problem with sequence dependent setups, GLSP: general lotsizing and scheduling problem, RH: rolling horizon, SCN: scenario approximation, PLA: piecewise linear approximation, RO: robust optimization, BLT: bivariate linearization technique, CDS: cumulative demand scenario, RS: pure random sampling, AS: antithetic sampling, LHS: Latin hypercube sampling, MM: moments matching, SRS: simple random sampling, DS: descriptive sampling.

As can be seen, the practical benefits of the static uncertainty strategy have attracted significant interest in the stochastic lot-sizing literature, although recently the focus has shifted towards the static-dynamic uncertainty strategy. In the model discussed in this chapter, we use the static uncertainty strategy. From the studies that do not assume dynamic demand, only Meistering and Stadtler (2017) consider seasonal demand in their numerical study, but this is done assuming that demand is time and product independent. In the context of the study reported in this thesis, it is important to note that although demand is usually assumed to be dynamic, the models and approximations developed in the studies reviewed can deal with demand seasonality, because this can be considered as a special case of dynamic demand. Moreover, the scenario representation of uncertainty remains widely used to solve stochastic lot-sizing problems. In most studies using a scenario representation of uncertainty, the scenario-generation method is unable to account for correlations. The studies able to account for this assume that demands of different product-period combinations are mutually independent. The only exception is Brandimarte (2006), in which demand is assumed to be cross-correlated, but not auto-correlated. Finally, we mention that the PLA approach is theoretically able to deal efficiently with demand dependencies, even though all studies utilizing this approximation strategy assume independent demands.

In this chapter, we analyze a SCLSP using a static uncertainty strategy in which demand can be static, dynamic or seasonal, and can be product and period-dependent. Unmet demand is assumed to be backordered and the amount of backorders is limited through service-level constraints. To approximate and solve the problem, we use a scenario approach.

4.2.2 Scenario generation

When a SCN approach is used to solve the SCLSP, it is crucial that the scenario-generation method creates scenarios that reflect all the important characteristics of uncertainty to obtain a solution of good quality. Specifically, in the case of interest, seasonal demand, the dependence of demand is such a characteristic. The majority of the SCLSP studies that use a SCN approximation to solve the problem, identified in the literature review in the previous section, use simple scenario-generation methods, such as SRS and DS, which are unable to account for demand dependencies. Alternatively, they use more complex methods, such as MM, that can deal with dependencies but are applied to simpler settings where correlations are disregarded. Of course, SP and scenario-generation techniques are not only applied to stochastic lot-sizing problems, or more generally to inventory management problems, but also to many other problems from a wide variety of fields. In the supply chain management area, SP is extensively used to solve supply chain network design problems. Other fields include finance, telecommunication and electricity and energy generation (Di Domenica et al., 2009). Because of the central role generally played by scenario-generation methods in stochastic problems and specifically in stochastic lot sizing, we provide a brief overview of the methods found in the scientific literature.

The most commonly used scenario-generation approach is sampling, which can take many forms (Kaut and Wallace, 2003). In its simplest but most prevalent form, sampling methods can sample from univariate distributions only, and therefore can handle multivariate distributions only if the random variables are independent. Monte Carlo sampling (e.g. the SRS approach mentioned earlier), importance sampling, DS and stratified sampling (e.g. LHS) are examples of this kind of technique. Other sampling methods can generate samples from multivariate distributions respecting possible correlations between the variables, albeit at the cost of significantly increasing complexity. Examples of such methods are found in Cario and Nelson (1997) and Deler and Nelson (2000), who use transformation-based approaches to obtain the desired multivariate random vector with an arbitrary correlation matrix. Another method that deals with correlations consists of conducting principal component analysis (PCA) on the original variables and then sampling scenarios of the principal components, which can then be transformed back into scenarios of the random variables that exhibit the correct correlations. As an example, Loretan (1997) uses PCA to develop stress scenarios for market risk in financial instruments. The sampling approaches mentioned so far assume that the dependencies between the random variables can be captured by standard correlations, whereas copulas-based methods do not assume this. An example of the use of copulas can be found in Kaut and Wallace (2011) and Kaut (2014), in which copulas are applied to single and multi-period portfolio optimization problems.

The main disadvantage of sampling-based methods is that they may generate scenario trees that have significantly different statistical properties than the sampled distribution with a limited number of scenarios (Mitra and Di Domenica, 2010). As a result, statistical methods were developed in order to obtain scenarios following the statistical properties of the random variables of interest. The most common method in this class is MM. This method was first presented in Hoyland and Wallace (2001), where it was applied to generate scenarios of returns of different financial assets with the goal of finding the optimal funds' allocation among those assets. This method, as opposed to many sampling methods, does not assume that the distribution functions of the marginal variables are known. MM consists of solving one or more nonlinear, non-convex optimization problems that which generate a predetermined number of scenarios which have statistical properties, including correlations, as close as possible to the ones of the sampled multivariate distribution. Other statistical approaches are path-based methods, which are employed when the underlying random process follows an econometric or time-series model. For example, Conejo et al. (2010) show how path-based methods can be used to generate unit-availability scenarios in the electricity markets. These methods are built on the assumption that demand is time-dependent and can accommodate cross-correlation, which, however, increases their complexity.

Finally, certain studies generate scenarios that combine two or more of the aforementioned approaches. For example, Schütz et al. (2009) analyze a SCND problem in the meat industry and generate demand scenarios using a combination of path-based methods, which take care of inter-temporal correlation, PCA, which takes care of cross-correlations between products, and finally MM, which ensures that the first four moments in the scenarios are as close as possible to those in the historical data. Calfa et al. (2014) present a similar approach that combines path-based methods and a modified version of MM that, apart from moments, also attempts to match the empirical cumulative distribution function of the data. The method is applied to forecast demand and production yield for illustrative purposes.

The above summary is not a comprehensive review of the literature available on the topic of scenario generation. However, we it is sufficient for the purpose of this thesis. For a more comprehensive review, the reader is referred to Mitra and Di Domenica (2010) and Kaut and Wallace (2003).

The reformulation of the scenario approximation of the SCLSP model that we propose requires scenarios of the cumulative demand for each product and period combination. These scenarios can be obtained independently, making it unnecessary in a practical setting to estimate correlations from limited or incomplete historical data and to consider them when generating scenarios, thus simplifying the latter process significantly. Moreover, our approach can be directly applied to problems in which demand dependencies are nonlinear, whereas the more sophisticated scenario-generation methods usually deal with linear dependencies (correlations) only. This might be useful in some settings when the demand of a product is seasonal. As an example, demands of different periods can be uncorrelated or negatively correlated in "normal" seasons, e.g. the timing of the demand during the season might change, but the total demand in the season remains unchanged. However, they might be highly positively correlated in very "good" or "bad" seasons, i.e. the demand in every period of the season soars or drops. In addition, in these settings, similar to the dependence over time, the dependence across products can also be nonlinear.

4.3 Problem statement and model formulation

We use the SCLSP with a static uncertainty strategy as a starting point for modeling the inventory management problem analyzed in this chapter. In line with the literature, and following, for the most part, the notation of Helber et al. (2013), we divide the (finite) planning horizon into T discrete time periods. We assume that there is a single facility, or resource, with a limited regular-time capacity of b_t, measured in time, in period t,

Table 4: SCLSP model – notation

	Sets
$\mathcal{K} = \{1, \dots, K\}$	Set of products
$\mathcal{T} = \{1, \dots, T\}$	Set of periods
	Deterministic parameters
b_t	Production capacity in period $t \in \mathcal{T}$
pt_k	Unit processing time of product $k \in \mathcal{K}$
sc_k	Setup cost of product $k \in \mathcal{K}$
st_k	Setup time of product $k \in \mathcal{K}$
h_k	Unit holding costs per period for product $k \in \mathcal{K}$
oc	Cost of one unit of overtime
	Random variables
D_{kt}	Demand of product $k \in \mathcal{K}$ in period $t \in \mathcal{T}$
NI_{kt}	Net inventory of product $k \in \mathcal{K}$ at the end of period $t \in \mathcal{T}$
IP_{kt}	Physical inventory of product $k \in \mathcal{K}$ at the end of period $t \in \mathcal{T}$
	Decision variables
x_{kt}	Lot size of product $k \in \mathcal{K}$ in period $t \in \mathcal{T}$
y_{kt}	Binary setup variable of product $k \in \mathcal{K}$ in period $t \in \mathcal{T}$
o_t	Overtime used in period $t \in \mathcal{T}$

which can produce all K products. The capacity can be expanded by the use of overtime, which comes at a cost of oc per unit of time. In order to produce product $k \in \mathcal{K}$ in a given period, a setup operation with corresponding costs sc_k and time st_k must be performed, and a processing time pt_k is required to produce each unit of the product. Holding costs per unit and period are h_k. Unmet demand is assumed to be backordered and we limit its amount using service-level constraints. The demand for product k in period $t \in \mathcal{T}$, D_{kt}, is a random variable with a known probability distribution, and dependencies between different product-period combinations may exist. Based on the notation introduced in Table 4, the stochastic capacitated lot-sizing model can be formulated as follows (Helber et al., 2013):

SCLSP Model

$$\text{Min } \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} h_k \cdot E[IP_{kt}] + \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} sc_k \cdot y_{kt} + \sum_{t \in \mathcal{T}} oc \cdot o_t \tag{40}$$

Subject to

$$NI_{kt} = NI_{k,t-1} + x_{kt} - D_{kt} \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T} \tag{41}$$

$$IP_{kt} = \max(0, NI_{kt}) \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T} \tag{42}$$

$$\sum_{k \in \mathcal{K}} (pt_k \cdot x_{kt} + st_k \cdot y_{kt}) \leq b_t + o_t \qquad \forall t \in \mathcal{T} \tag{43}$$

$$x_{kt} \leq \frac{b_t}{pt_k} \cdot y_{kt} \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T} \tag{44}$$

$$\text{Service level constraints} \tag{45}$$

$$x_{kt} \geq 0 \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T} \tag{46}$$

$$o_t \geq 0 \qquad \forall t \in \mathcal{T} \tag{47}$$

$$y_{kt} \in \{0, 1\} \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T}. \tag{48}$$

The objective (40) is to minimize expected total costs, which are composed of expected inventory holding costs, setup costs and overtime costs. Given the uncertainty of demand, the physical inventory is itself a random variable, which is defined in the constraints. Constraints (41) define the stochastic net inventory, which is a function of the net inventory in the previous period, the produced quantity and the random demand. Constraints (42) define the random physical inventory, which is a function of the net inventory. Constraints (43) ensure that the capacity used for producing and performing setups does not exceed the sum of the regular and overtime capacities. Constraints (44) ensure that in order to produce an item, the required setup is performed. Constraints (45) are service-level constraints which limit the amount of backorders; their exact form depends on the service-level measure used and will be discussed shortly. Constraints (46)-(48) are variable definition constraints. As previously mentioned, the static uncertainty strategy applied here can be used with any other production model, e.g. the GLSP or CLSD, by simply substituting the production-related constraints (43), (44), (46) and (48) by the (deterministic) production constraints of the relevant model.

In the setting with uncertain seasonal demand, analyzed in this chapter, inventory plays three roles. First, it allows exploiting economies of scale by producing in large batches, which creates inventory, and thus save on setup costs. The inventory kept for this purpose is called cycle stock. Second, it serves as a buffer against demand uncertainty to ensure that the desired service-level target is met; this inventory is referred to as safety stock. Third, inventory enables to serve demand in the peak season by pre-producing in the off-season instead of using the expensive overtime capacity in the peak season; this inventory is named seasonal stock. Although in the SCLSP model they are not explicitly differentiated, all of these different inventory types are included in the optimal inventory plan.

At this point it is important to discuss the topic of unmet demand and the importance of limiting its amount. Whenever a firm is unable to meet a customer's demand, this demand can either be lost or backordered, i.e. served at a later point in time, possibly at a discount. Accordingly, this should be reflected in inventory models, which therefore are commonly developed under the assumption that unmet demand is either lost or backordered. Of course, in practical settings, a mix of these two assumptions is likely to be the most accurate representation of reality. As mentioned in the literature review section, the common assumption in inventory theory is the one of backordering, which has the advantage of simplifying the model formulation. However, it should be noted that when the backordered quantity is kept low enough, the difference between the model formulations under the two assumptions is negligible (Azoury and Miyaoka, 2013). For these reasons, we also assume that unmet demand is backordered in our model. Subsequently, the amount of backorders must be controlled.

In inventory models, and specifically in lot-sizing problems, as shown in the literature review, backorders are limited by either attaching a cost to each backordered unit or by enforcing a certain minimum service level, which measures the degree of customer orders' satisfaction according to some metric. Due to the practical difficulty of quantifying backorder costs, we use the latter approach to limit the size of backorders. Because several service-level measures have been developed in the inventory management literature, in the following only the most common ones will be presented and formulated for the SCLSP model.

The α service level measures the probability that no stockout occurs during a production cycle, i.e. the time between two consecutive production periods, or in a given time interval. The α service level of a product k for a certain period t is thus defined as

$$\alpha_{kt} = Pr(NI_{kt} \geq 0). \tag{49}$$

If the α service level measure is used to control backlogs, then, given a target service level per product and period α_k, constraints (45) would take the following form (Koca et al., 2015):

$$Pr(NI_{kt} \geq 0) \geq \alpha_k \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T}. \tag{50}$$

This measure is intuitive and mathematically tractable, but it has the downside that it only measures the probability of the occurrence of a stockout without considering its magnitude.

The β service level, or fill-rate, measures the percentage of demand in a given time span that is met immediately, that is, without backordering. The time span considered can be either any finite number of periods, or a production cycle. In this thesis, we use the terminology employed by Helber et al. (2013): the *backorder* in period t refers to the amount of products demanded in period t which is not met in the same period, whereas the *backlog* in period t refers to the quantity of backorders generated in periods 1 to t that remain unfilled at the end of period t. For a given product k, we define the number of units (newly) backordered in period t as B_{kt}, i.e. the amount of demand of period t which is not met in the same period. The so-called *finite horizon* fill-rate of product k, β_{kT}, can then be defined as follows:

$$\beta_{kT} = 1 - \frac{E[\sum_{t=1}^{T} B_{kt}]}{E\big[\sum_{t=1}^{T} D_{kt}\big]}. \tag{51}$$

The backorders in period t can be obtained by subtracting the backlog at the beginning of period t, i.e. immediately after production but before the period's demand is satisfied, from the backlog at the end of period t (Tempelmeier, 2013). We can therefore define B_{kt} as

$$B_{kt} = \max(0, -NI_{kt}) - \max\Big(0, \sum_{\tau=1}^{t-1} D_{k\tau} - \sum_{\tau=1}^{t} x_{k\tau}\Big). \tag{52}$$

To implement a finite-horizon β service-level constraint (where the horizon is the entire planning horizon) with product-specific targets β_{kT}, constraints (45) need to be replaced with:

$$B_{kt} = \max(0, -NI_{kt}) - \max\Big(0, \sum_{\tau=1}^{t-1} D_{k\tau} - \sum_{\tau=1}^{t} x_{k\tau}\Big) \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T} \tag{53}$$

$$1 - \frac{\sum_{t\in\mathcal{T}} E[B_{kt}]}{\sum_{t\in\mathcal{T}} E[D_{kt}]} \geq \beta_{kT} \qquad \forall k \in \mathcal{K}. \tag{54}$$

In order to implement a *cycle* β service-level constraint in the SCLSP, to measure the percentage of demand that is met immediately in a given production cycle, more changes are necessary, because there is a need to track every cycle of every product in the planning horizon considered. First, for a given product k and for two given consecutive order cycles $(i-1)$ and i, ending in periods τ_{i-1} and τ_i, respectively, we can define the β cycle service level for cycle i of product k, $\beta_k(\tau_i)$, as

$$\beta_k(\tau_i) = 1 - \frac{E\big[\sum_{t=\tau_{i-1}}^{\tau_i} B_{kt}\big]}{E\big[\sum_{t=\tau_{i-1}}^{\tau_i} D_{kt}\big]}. \tag{55}$$

The numerator on the right-hand side of equation (55) represents the new backorders generated during cycle i and the denominator represents the total demand in that cycle, thus defining the cycle service level as intended. Second, we must define two new tracking variables: l_{kt}, which counts the number of periods since the last setup of product k prior to t, and w_{kt}, which equals one if a setup for product k is performed in period $t+1$ and zero otherwise.

Finally, denoting the β cycle service-level target for each cycle of product k by β_k, we can substitute constraints (45) with the following constraints (Tempelmeier, 2013):

$$B_{kt} = \max(0, -NI_{kt}) - \max(0, \sum_{\tau=1}^{t-1} D_{k\tau} - \sum_{\tau=1}^{t} x_{k\tau}) \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T} \qquad (56)$$

$$l_{kt} = (l_{k,t-1} + 1) \cdot (1 - y_{kt}) \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T} \qquad (57)$$

$$l_{k0} = -1 \qquad \forall k \in \mathcal{K} \qquad (58)$$

$$w_{kt} = y_{k,t+1} \qquad \forall k \in \mathcal{K},\ t = 1, \ldots, T-1 \qquad (59)$$

$$w_{kT} = 1 \qquad \forall k \in \mathcal{K} \qquad (60)$$

$$1 - \frac{E\big[\sum_{\tau=t-l_{kt}}^{t} B_{k\tau}\big]}{E\big[\sum_{\tau=t-l_{kt}}^{t} D_{k\tau}\big]} \geq \beta_k \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T} | w_{kt} = 1. \qquad (61)$$

Constraints (56) define the backorders in period t, as already explained. Constraints (57) and (58) are used to measure the lengths of cycles. For every time period t, variable l_{kt} is equal to zero if the current production cycle started in period t or equal to the length of the production cycle if it started in a previous period. Constraints (59) and (60) are used to define the end of the current production cycle of product k. For every time period t, variable w_{kt} is equal to one if the current production cycle of product k will terminate in period $t+1$ or if t is the last period of the planning horizon, and equal to zero otherwise. Service-level constraints (61) ensure that for each period t, if a production cycle of product k started in period $t - l_{kt}$ ends, then the cycle service-level target for that cycle is met. Tempelmeier et al. (2018) show a linear reformulation of constraints (56)-(61), which does not require the additional variables l_{kt} and w_{kt}, obtained by expressing the cycle service level in terms of an equivalent finite-horizon fill-rate. This reformulation ensures that if capacity constraints are not binding, then the target cycle service level is met exactly, and otherwise it guarantees that the target is not missed. As noted by Helber et al. (2013), apart from the difficulty of simultaneously making production quantity decisions and enforcing service-level constraints, the β service-level measure also has the disadvantage of only controlling the size of the backorders in a production cycle without reflecting the time that is necessary to meet this backordered demand.

The γ service level attempts to overcome this shortcoming of the fill-rate measure by reflecting the waiting time of customers whose demand is backordered, to a certain extent. Denoting the existing backlog of product k at the end of period t by BL_{kt}, the γ service level of product k in period t can be defined as

$$\gamma_{kt} = 1 - \frac{E[BL_{kt}]}{E[D_{kt}]}. \qquad (62)$$

Alternatively, the measure can be defined as an average over the whole planning horizon. The problem with this measure is that for some periods it might be either negative, if the expected backlog is larger than the expected demand in that period, or even undefined, if the expected demand is zero. Moreover, this measure is not clearly interpretable.

The δ service level is a measure introduced by Helber et al. (2013) which simultaneously measures the size of the backorders and the waiting time experienced by customers. Moreover, it is well-defined and offers a clear interpretation. Averaged over the entire planning horizon, the δ service level of product k is defined as

$$\delta_k = 1 - \frac{\sum_{t=1}^{T} E[BL_{kt}]}{\sum_{t=1}^{T}(T-t+1) \cdot E[D_{kt}]}. \tag{63}$$

To interpret this measure, we can use Little's Law. We can define the customers waiting for their backorders of product k as the inventory of this production system. The average inventory of the system is then $\frac{1}{T} \cdot \sum_{t=1}^{T} E[BL_{kt}]$. If long-run production matches long-run demand, then the system has an average throughput of $\frac{1}{T} \cdot \sum_{t=1}^{T} E[D_{kt}]$. The average waiting time for customers to receive their backlogged demand, $E[W_k]$, is therefore

$$E[W_k] = \frac{\frac{1}{T} \cdot \sum_{t=1}^{T} E[BL_{kt}]}{\frac{1}{T} \cdot \sum_{t=1}^{T} E[D_{kt}]}. \tag{64}$$

Using this equality, we can obtain another, equivalent, definition of the δ service level:

$$\delta_k = 1 - \frac{\sum_{t=1}^{T} E[D_{kt}]}{\sum_{t=1}^{T}(T-t+1) \cdot E[D_{kt}]} \cdot E[W_k]. \tag{65}$$

The second term on the right-hand side of this equality is the percentage of the maximum demand-weighted waiting time that customers of product k experience. Indeed, the denominator of this term is the maximum demand-weighted waiting time, because every period's demand is multiplied with the maximum waiting time for that demand, i.e. the time left until the end of the planning horizon. The numerator measures the actual demand-weighted waiting time experienced by customers, because each period's demand is multiplied by the average waiting time. Therefore, we can conclude that the δ service level measures "the expected percentage of the maximum possible demand-weighted waiting time that the customers of product k are protected against" (Helber et al., 2013). This measure reflects both the size of backorders and the waiting time experienced by customers, and is clearly well-defined. Indeed, if the entire demand is always met immediately, i.e. there are no backlogs in any period, or, equivalently, the average waiting time is zero, then $\delta_k = 1$. If, instead, the total production quantity in the planning horizon is zero, meaning that the total backlog and the average waiting time are at their respective maximum values, then $\delta_k = 0$. All other possible cases are somewhere between these two extreme events, therefore $0 \leq \delta_k \leq 1$.

Denoting the service-level target of product k by δ_k, we can implement a δ service-level constraint by substituting constraints (45) with the following (Helber et al., 2013)

$$BL_{kt} = \max(0, -NI_{kt}) \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T} \tag{66}$$

$$1 - \frac{\sum_{t=1}^{T} E[BL_{kt}]}{\sum_{t=1}^{T} (T - t + 1) \cdot E[D_{kt}]} \geq \delta_k \qquad \forall k \in \mathcal{K} \tag{67}$$

$$\sum_{t \in \mathcal{T}} x_{kt} \geq \sum_{t \in \mathcal{T}} E[D_{kt}] \qquad \forall k \in \mathcal{K}. \tag{68}$$

In the SCLSP analyzed in this chapter, the δ service level is chosen to limit backorders, given the above-mentioned properties of this measure.

The SCLSP with δ service-level constraints is a mixed integer nonlinear programming model (MINLP), with the nonlinearity originating from the expected inventory and expected backlog functions, which are defined as the expectation of maximum functions. Because of its nonlinear nature, the most common approach to solving the SCLSP is to approximate it by linearizing the nonlinear expected inventory and backlog functions, either with the use of scenarios or of piecewise linear functions. We now analyze the scenario approximation of the problem, focusing on the setting with dependent demands.

4.4 A cumulative scenario approximation of the SCLSP

4.4.1 The standard scenario approximation of the SCLSP

As mentioned, a common approach in the stochastic lot-sizing literature (and in SP in general) is to represent the uncertainty in the form of scenarios, i.e. possible discrete realizations of the stochastic parameters. It is assumed that one scenario s of a set $\mathcal{S} = \{1, \ldots, S\}$ of scenarios will be realized, and the ex-ante probability that this will happen is pr_s. These scenarios are paths of demand realizations used to approximate the total costs and expected service level of a static uncertainty production plan. This discretization of the demand process transforms the SCLSP model into a linear model, because the expected value of the net inventory, and therefore of the physical inventory and backlog, can be expressed as linear functions, specifically as weighted sums of the values of the variable in each scenario, where the weights are the probabilities of the scenarios occurring.

Using the SP terminology, the resulting problem is a two-stage stochastic program. All production decisions are first stage decision variables, i.e. scenario independent, because they must be taken before the uncertainty unfolds. Net inventory, physical inventory and backlogs are second stage variables, i.e. scenario dependent, given that they depend on the realization of the demand uncertainty.

Table 5: SCLSP-SCN model – additional/modified notation

	Sets
$\mathcal{S} = \{1, \ldots, S\}$	Set of scenarios
	Parameters
pr_s	probability of occurrence of scenario $s \in \mathcal{S}$
D_{kts}	Demand of product $k \in \mathcal{K}$ in period $t \in \mathcal{T}$ in scenario $s \in \mathcal{S}$
$E[D_{kt}]$	Expected demand of product $k \in \mathcal{K}$ in period $t \in \mathcal{T}$
δ_k	Delta service-level target for product $k \in \mathcal{K}$
	Scenario-dependent decision variable
NI_{kts}	Net inventory of product $k \in \mathcal{K}$ at the end of period $t \in \mathcal{T}$ in scenario $s \in \mathcal{S}$
IP_{kts}	Physical inventory of product $k \in \mathcal{K}$ at the end of period $t \in \mathcal{T}$ in scenario $s \in \mathcal{S}$
BL_{kts}	Backlog of product $k \in \mathcal{K}$ at the end of period $t \in \mathcal{T}$ in scenario $s \in \mathcal{S}$

Using the notation shown in Tables 4 and 5, and indicating that $E[D_{kt}] = \sum_{s \in \mathcal{S}} D_{kts}$, the scenario-approximation model SCLSP-SCN with δ service-level constraints is formulated as (Helber et al., 2013):

SCLSP-SCN Model

$$\text{Min} \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} h_k \cdot \Big(\sum_{s \in \mathcal{S}} pr_s \cdot IP_{kts} \Big) + \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} sc_k \cdot y_{kt} + \sum_{t \in \mathcal{T}} oc \cdot o_t \tag{69}$$

Subject to constraints (43), (44),(46)-(48) and

$$NI_{kts} = NI_{k,t-1,s} + x_{kt} - D_{kts} \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \tag{70}$$

$$IP_{kts} \geq NI_{kts} \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \tag{71}$$

$$BL_{kts} \geq -NI_{kts} \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T},\ \forall s \in \mathcal{S} \tag{72}$$

$$\sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{S}} pr_s \cdot BL_{kts} \leq (1 - \delta_k) \cdot \sum_{t \in \mathcal{T}} \big((T - t + 1) \cdot E[D_{kt}] \big) \qquad \forall k \in \mathcal{K} \tag{73}$$

$$\sum_{t\in\mathcal{T}} x_{kt} \geq \sum_{t\in\mathcal{T}} E[D_{kt}] \qquad \forall k \in \mathcal{K} \tag{74}$$

$$IP_{kts}, BL_{kts} \geq 0 \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T},\ \forall s \in \mathcal{S}. \tag{75}$$

Constraints (70) calculate the net inventory in each scenario and replace constraints (41). Constraints (71) and (72), together with constraints (75), define physical inventory and backlogs for all scenarios, respectively; the δ service-level constraints (73) and the direction of the optimization ensure that both IP_{kts} and BL_{kts} are well-defined using these constraints. As explained, in objective function (69) and constraints (73), the nonlinear expected inventory and backlog functions are calculated as the probability-weighted average of inventory and backlogs in all scenarios, respectively. Constraints (74) ensure that the inventory system is stable.

In Section 4.2.1, a short overview of the scenario-generation literature was presented. We now discuss certain frequently used scenario-generation methods in more detail.

As noted, many stochastic lot-sizing studies assume that the demand for each product-period combination is independent from all other combinations. As a result of this assumption, sampling is simple, because samples for each product-period couple can be drawn independently. In the majority of these cases, the sampling method used is SRS. Using this method, to obtain a scenario s, a (pseudo-) random number is first drawn from the standard uniform distribution in the interval $[0, 1]$ for each product-period couple (k, t) independently, denoted by v_{kts}; this number can be interpreted as a cumulative probability. Second, using the inverse cumulative distribution function of the relevant demand, $F^{-1}_{D_{kt}}$, a random value of the demand is obtained. Specifically, the demand for product k in period t in scenario s is

$$D_{kts} = F^{-1}_{D_{kts}}(v_{kts}). \tag{76}$$

This procedure is then repeated S times, where S is the number of desired scenarios, each of which has a probability of occurrence equal to $1/S$. When the number of scenarios is limited, as in most SP applications, the resulting demand sample's cumulative distribution function might be considerably different from the corresponding sampled distribution function. Additionally, the optimal objective function and optimal solution of the model is likely to vary considerably from one sample to another.

To overcome these issues, DS was proposed by Saliby (1990). Given a predetermined number of target scenarios S, the first step of this method is to deterministically obtain, for each product-period combination, S demand scenarios. To do this, S equidistant numbers are taken from the interval $[0, 1]$, representing quantiles of the sampled distribution. Next, the demand values corresponding to the quantiles are obtained using the relevant inverse cumulative distribution function. Specifically, the deterministic set of values of the uncertain demand is obtained by applying the following formula:

$$D_{kts} = F^{-1}_{D_{kt}}\Big(\frac{s-0.5}{S}\Big) \qquad\qquad s = 1, \dots, S. \tag{77}$$

Assigning a probability of $1/S$ to each demand realization, the cumulative distribution function of this deterministic sample of demand is expected to be a reliable representation of the underlying sampled distribution function. In a second step, for each couple (k, t), the elements of this deterministic sample are "shuffled", i.e. the sampled values are randomly assigned to specific scenario numbers $s = 1, \dots, S$, creating a random permutation of the deterministic elements of the set. Now, a given scenario $s \in \{1, \dots, S\}$ represents a path of demand from the beginning to the end of the planning horizon, and D_{kts} represents the demand of product k in period t in that specific scenario. In addition, each scenario has the same probability of occurrence equal to $1/S$. Obtaining different sets of scenarios is possible by re-shuffling the deterministic demand values for each period-product combination. This leads to a lower variance of the objective function and optimal solution than SRS.

A method that can be thought of as a mix of DS and SRS is univariate LHS. Using this method, for each product-period combination, the interval $[0, 1]$ is divided into S separate portions, all equal in size, and then a random value v_{kts} is drawn from each portion $s = 1, \dots, S$. The value of demand of product k in period t corresponding to v_{kts} is then obtained again by using the inverse cumulative distribution function as follows:

$$D_{kts} = F^{-1}_{D_{kt}}(v_{kts}) \qquad\qquad s = 1, \dots, S. \tag{78}$$

To obtain the S scenarios of demand, these values are then shuffled using the same technique as used with DS. Now, for a given $s \in \{1, \dots, S\}$, D_{kts} represents the demand of product k in period t in scenario s. The sampling cumulative distribution function in this case would also be quite close to the sampled one.

Out of the statistical scenario-generation methods mentioned in Section 4.2.2, MM is the most commonly used. To apply it to the static uncertainty SCLSP, the number of scenarios S of demand and the Q statistical properties to be matched by them must be determined in advance. Denote by d the vector of demand of all products, periods and scenarios combinations, and by pr the vector of the probabilities of the S scenarios, where the s^{th} element, pr_s, corresponds to the probability of scenario s. Also, denote the weight of statistical property q by w_q, the mathematical expression of statistical property q by $f_q(\cdot,\cdot)$ and the value of the statistical property q to be matched by $Sval_q$. The scenario-generation optimization model can be formulated as follows (Calfa et al., 2014):

$$\min_{d,p} \sum_{q=1}^{Q} w_q \cdot (f_q(d,pr) - Sval_q)^2 \tag{79}$$

Subject to

$$\sum_{s=1}^{S} pr_s = 1 \tag{80}$$

$$pr_s \in [0,1] \qquad \forall s = 1,\ldots,S. \tag{81}$$

$$\tag{82}$$

The objective function (79) minimizes the weighted sum of the (squared) deviations of the values of the statistical properties of the scenario set from their target values. Examples of these statistical properties are the first four moments of the marginal distribution of each product-period combination. The weights enable the modeler to choose the relative importance of meeting the different statistical properties, in case not all of them can be met perfectly. Constraint (80) ensures that the probabilities of the scenarios, which are the decision variables of this optimization problem together with the values of the random variables for each scenario, sum up to one. Finally, constraints (81) make sure that the probabilities of the scenarios are larger than zero, and thus well-defined. Although in the stochastic lot-sizing literature this method has been applied assuming no demand correlations, it can be applied to a setting with both auto- and cross-correlated demand because the auto- and cross-correlation of demand can be part of the statistical properties to be matched by the scenarios. For example, to ensure that the scenarios exhibit the correct cross-correlation in all periods, one set of the $f_q(\cdot,\cdot)$ functions would be defined as:

$$corr_{k,j} = \sum_{s=1}^{S} (D_{kts} - mean_{kt}) \cdot (D_{jts} - mean_{jt}) \cdot pr_s \quad \forall (k,j) \in \mathcal{K} | k \neq j, \ \forall t \in \mathcal{T}, \tag{83}$$

where $mean_{kt}$ is the mean of the demand of product k in period t. The deviation of $corr_{k,j}$ from its target value would then be entered into the objective function (79).

If the objective function value is small (the best possible outcome being zero) then the solution scenario set satisfies the desired statistical properties sufficiently well. The MM model is in general nonlinear and non-convex, and thus the solution is likely to be a local optimum (Hoyland and Wallace, 2001). It is possible that multiple scenario sets match the statistical properties as intended, however, they might significantly differ from the original data distribution (Mitra and Di Domenica, 2010).

Other common scenario-generation methods that are used when demand is time-dependent are path-based methods. Demand is assumed to follow an econometric or time-series model, which is used to derive scenarios. Conejo et al. (2010) show how to develop scenarios following one of these methods. An auto-regressive integrated moving average (ARIMA) model is fitted to the demand of each product k. The demand process can thus be expressed as

$$\left(1-\sum_{j=1}^{p}\phi_j\cdot B^j\right)\cdot(1-B)^d D_{kt}=\left(1-\sum_{j=1}^{q}\theta_j\cdot B^j\right)\cdot\varepsilon_{kt}, \tag{84}$$

where B is the so-called lag operator, i.e. $B^j \cdot D_{kt} = D_{k,t-j}$. The general model in (84) has p autoregressive parameters $\phi_1, \ldots, \phi_p$, q moving average parameters $\theta_1, \ldots, \theta_q$, and a degree of differencing d. The errors ε_{kt} are generally assumed to be independent and identically distributed (iid) and to follow a normal distribution. Although the standard model in (84) is not directly applicable to a seasonal demand series, there is an extension to this model that can be applied to these cases.

Once the parameters of the ARIMA model are determined, scenarios are obtained following an iterative procedure which generates one scenario path per iteration. To generate one scenario path for product k, for each period t a random sample from the distribution of ε_t is drawn and the ARIMA model (84) is used to calculate D_{kts} based on the demand and random errors obtained in the periods up to t in scenario s. This basic scenario-generation procedure must be modified when there are correlations among the error terms of different products, because the paths cannot be sampled for each product individually. Conejo et al. (2010) present a procedure to do that, based on an orthogonal transformation of the normal errors generated for each product independently.

To conclude, in the presence of correlations, the sampling methods outlined above cannot be directly applied. This is because it is impossible to draw samples from the marginal demand distribution of each (k, t) pair independently and then combine different realizations of different pairs randomly to create scenarios. Instead, samples must be drawn from a joint distribution with correlated random variables, which is a challenging task. Both the MM and path-based methods can deal with linear dependencies, but only at the cost of an increased complexity in their approaches. Moreover, it is also a challenge to estimate the dependencies between the demand of different (k, t) pairs in the first place, particularly in practical applications where data availability is an issue. This is especially likely to be the case when demand is seasonal, because every seasonal cycle provides only a single data point for the estimation of the demand distribution.

4.4.2 The new cumulative scenario approximation of the SCLSP

We now propose an alternative formulation of the scenario approximation of the SCLSP, which significantly simplifies the scenario-generation process. The simplification is due to a different definition of scenarios, that eliminates the necessity to consider demand dependencies when generating scenarios and, therefore, to estimate them in the first place. We first describe how to obtain the new scenario approximation and then compare it in detail to the standard scenario approximation used in the literature, clearly pointing out the advantages of our new approach.

To justify using this approach, we note that the net inventory, and therefore the physical inventory and the backlog, of product k at the end of period t is a function of the *cumulative* demand up to period t. This means that in order to determine the net inventory, physical inventory, and backlog of the product at the end of a period, we do not need to know the demands for the product in the individual periods up to that period, but only their sum. This is the key to our approach.

In the SCLSP model, we must calculate the expected value of the total inventory costs and of the total backlogs for each product. The SCLSP-SCN approach essentially calculates these values as $E\big[\sum_{t\in\mathcal{T}}\sum_{k\in\mathcal{K}} h_k \cdot IP_{kt}\big]$ and $E\big[\sum_{t\in\mathcal{T}} BL_{kt}\big]$, respectively. The estimates are obtained with the use of scenarios generated from the joint probability distribution of demands for all products in all periods. We propose an alternative procedure for obtaining an approximation of the two expected values of interest, based on the fact that the expectation and summation operators can be interchanged. That is, we calculate the expected total inventory costs and total backlogs as $\sum_{t\in\mathcal{T}}\sum_{k\in\mathcal{K}} E[h_k \cdot IP_{kt}]$ and $\sum_{t\in\mathcal{T}} E[BL_{kt}]$, respectively. We then exploit the fact that both IP_{kt} and BL_{kt} depend on the cumulative demand for product k until period t. Hence, we estimate the corresponding expected values by using scenarios generated from the probability distribution of the corresponding cumulative demand. This new approach employs a cumulative demand "view" and we name the resulting approximation model *SCLSP-CDS*, where CDS stands for cumulative demand scenario.

To formulate the SCLSP-CDS model, we start by considering the original nonlinear SCLSP model (40)-(48), and denoting by CD_{kt} the random cumulative demand of product k up to period t, i.e. $CD_{kt} = \sum_{\tau=1}^{t} D_{k\tau}$. We can rewrite constraints (41) of the model as

$$NI_{kt} = \sum_{\tau=1}^{t} x_{k\tau} - CD_{kt} \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T}. \tag{85}$$

Therefore, adapting the SCLSP-SCN model to the cumulative demand view requires only two changes. First, constraints (70) are reformulated by expressing the net inventory of product k at the end of period t in scenario s as a function of the cumulative demand in scenario s, CD_{kts}, using the logic of constraints (85). Second, the δ service-level constraints (73) and constraints (74) are rewritten by using scenarios of cumulative demand to express the right-hand side of the inequalities. The SCLSP-CDS model therefore has objective (69), subject to constraints (43), (44),(46)-(48), (71), (72), (75) and the following three constraints:

$$NI_{kts} = \sum_{\tau=1}^{t} x_{k\tau} - CD_{kts} \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T},\ \forall s \in \mathcal{S}. \tag{86}$$

$$\sum_{t\in\mathcal{T}} \sum_{s\in\mathcal{S}} pr_s \cdot BL_{kts} \leq (1-\delta_k) \cdot \sum_{t\in\mathcal{T}} \sum_{s\in\mathcal{S}} \big(CD_{kts} \cdot pr_s\big) \qquad \forall k \in \mathcal{K} \tag{87}$$

$$\sum_{t\in\mathcal{T}} x_{kt} \geq \sum_{s\in\mathcal{S}} CD_{kTs} \qquad \forall k \in \mathcal{K}. \tag{88}$$

Compared to the SCLSP-SCN model, scenarios are not "paths" of demand in a scenario tree, instead they are scenarios of cumulative demands. In other words, the set $\mathcal{S}$ no longer contains S scenarios of the demand realization for all products for the entire planning horizon, instead it contains S scenarios of cumulative demand for each product-period combination. In principle, the number of scenarios could vary from one product-period combination to another, in which case we would define $K \times T$ sets of scenarios, defined as $\mathcal{S}_{kt} = \{1, \dots, S_{kt}\}$.

As for the standard SCLSP-SCN model, a scenario-generation method is needed to obtain the cumulative demand scenarios. The main advantages of the SCLSP-CDS lie exactly in this step, because it is unnecessary to account for demand dependencies when generating scenarios. This is clearly explained in the next section, where a more in-depth comparison between the SCLSP-SCN and SCLSP-CDS models is presented.

4.4.3 Comparing SCLSP-SCN and SCLSP-CDS

Figure 24 helps to explain the difference in the logic of estimating the total expected inventory in the SCLSP-SCN model and the SCLSP-CDS model. A similar figure and discussion can be used to explain the difference in the estimation of backlogs, but we focus on inventory here for ease of exposition. A simple setting with $K = T = 2$ and $S = 4$ is analyzed for illustration purposes. The following description of the figure clarifies the simplified notation used in the graphs.

The logic behind the SCN model is shown in Figure 24a. Scenarios are generated as paths through a so-called scenario fan, a special form of a scenario tree. The bold arrows, for example, characterize scenario (path) $s = 1$. Given a plan, the resulting total inventory costs in the scenario is calculated and denoted by $TIC_{s=1}$. The same is done for the other three scenarios shown in the figure. Finally, the expected total inventory costs, $E[TIC]$,

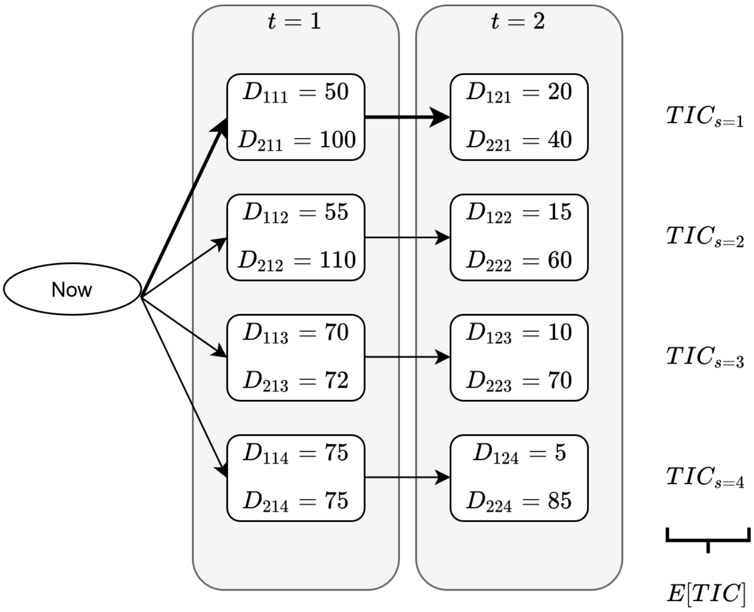

(a) SCLSP-SCN logic

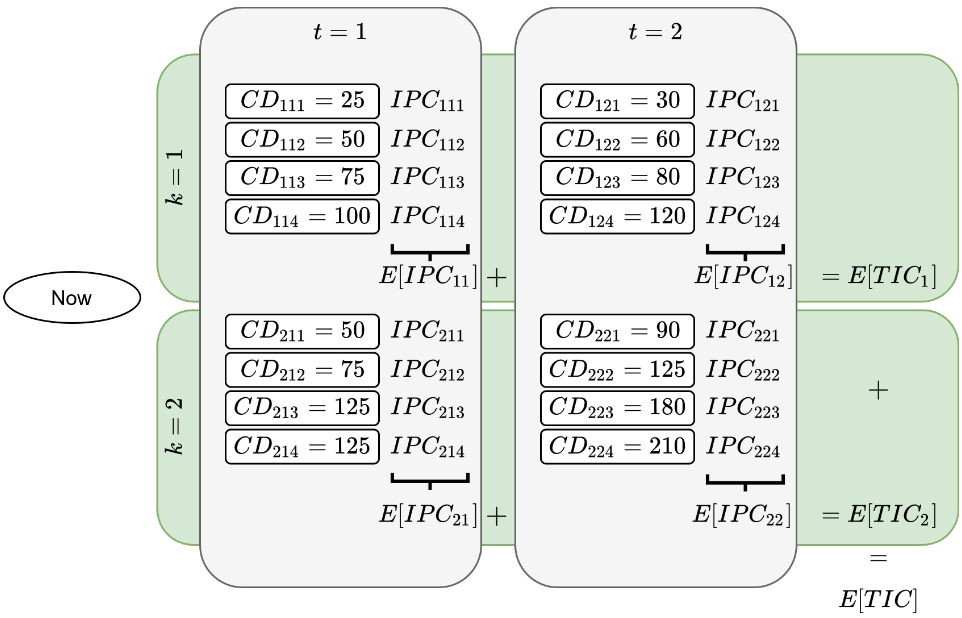

(b) SCLSP-CDS logic

Figure 24: Difference between SCLSP-SCN and SCLSP-CDS

are calculated as the probability-weighted sum of the total inventory of the four scenarios. As a comparison, the logic behind the CDS model is shown in Figure 24b. The estimation of the total inventory costs is decomposed into the estimation of single and independent inventory costs for each (k, t) pair: the two green rounded rectangles decompose the estimation problem into single-product problems, and the two grey rounded rectangles further decompose the problem into single-period problems (note that there are no arrows connecting the different (k, t) pairs, in contrast to the standard scenario fan). Scenarios for the cumulative demand are generated for each pair. To each such scenario corresponds a certain inventory costs figure, IPC_{kts}. The expected inventory costs for each (k, t) pair, $E[IPC_{kt}]$, is then calculated as the probability-weighted sum of inventory costs over all scenarios. Then, the expected total inventory costs of each product over the planning horizon, $E[TIC_k]$, is calculated by adding its expected inventory costs over all periods. Finally, these values are added for all the products to obtain the expected total inventory costs, $E[TIC]$.

The cumulative view has two clear advantages over the standard view when it comes to the scenario-generation process, due to estimating the expected inventory costs (and backlogs) for each (k, t) pair separately and then adding these estimates up. First, it is unnecessary to consider demand dependencies between products or periods when generating scenarios of cumulative demands. Indeed, the decomposition of total inventory into the sum of total inventory over all products (the green rectangles in figure 24b) eliminates the need for controlling the demand dependence between products in the scenarios. The further decomposition of each product's total inventory into the sum of inventory over all periods (the grey rectangles in figure 24b) eliminates the need for controlling the time-dependence of demand in the scenarios, because sampling from the cumulative demand distributions obviously ensures that the time-dependence is respected in the (implicit) path to arrive at any cumulative demand value. We emphasize that this decomposition only relates to the problem of estimating the expected inventory and backlog functions of different products, and not the production planning problem. The first stage production variables of different products are still linked through the capacity constraints, however, the evaluation of the resulting production plan in terms of inventory and backlog in the second stage can be performed independently for each product. Second, as a result of this simplification, it is unnecessary to *estimate* the dependencies between all product-period pairs in the first place. As mentioned, this is clearly an advantage in practice, because estimating the dependencies in problems with large amounts of products and periods is a challenging task, especially with limited data, which is likely an issue in settings with seasonal demand patterns.

The sampling methods presented in section 4.4.1 can be applied to derive the cumulative demand scenarios in the same way as in the standard SCN model, the only difference is that the sampled distributions are now the cumulative demand distributions instead of single period demand distributions. Although the SRS procedure is essentially unchanged, the application of DS and univariate LHS is slightly different. Specifically, in both scenario-generation methods, the shuffling procedure is no longer needed, because of the independence of the estimates of the inventories and backlogs of different (k, t) pairs. This means that DS becomes completely deterministic, whereas LHS still induces some variability in the estimates of the inventories and backlogs from one scenario set to another, although this variability is much smaller than in the standard SCN model. In addition, certain statistical scenario-generation methods, e.g. the MM and path-based methods shown in Section 4.4.1, can still be applied as before. However, the demand dependencies no longer have to be considered as part of the statistical properties to match. Because of the superiority of DS over SRS in the standard SCN approximation of the SCLSP, proven by Helber et al. (2013), we chose to use DS to derive the scenarios for the SCLSP-CDS model.

For completeness, we to point out that the PLA approach used to solve the SCLSP (see Helber et al., 2013 for details), which is not our research focus in this chapter, also employs the cumulative demand view underlying the CDS model. Indeed, in this alternative approximation approach, the expected backlog and inventory functions of different product-period combinations are approximated independently using piecewise linear functions and cumulative demand distributions.

We conclude mentioning that the cumulative view approach also has some disadvantages. First, we cannot evaluate the performance of a given plan for a specific scenario, because this requires that a concrete scenario path is constructed. This means that correlations should first be estimated and then enforced when generating the scenario, and it is precisely this step that the cumulative approach eliminates. In other words, without creating scenario paths we can only determine the expected performance of a plan, but have no information concerning its variance. This could be important if the decision maker wants to control risk. Second, this approach works only for a static uncertainty strategy (or for a single-stage SP in general). If a different strategy is used, we must make decisions at a stage based on information concerning the actual realization of demand before that stage and the expectation of future demand realizations that are conditional on the scenario path up to that stage. Third, it is impossible to measure the newly generate backorders in any given period using linear constraints, thus making it impossible to control backorders using popular service-level measures other than the δ one, e.g. the finite-horizon and cycle β service level measures.

4.5 Numerical study

The main advantage of the CDS approach is the simplification of the estimation and scenario-generation processes in settings with correlated demands. Because both processes are performed before solving the SCLSP, quantifying this advantage over other scenario-generation methods in a computational study is not a trivial matter. Therefore, in the numerical study presented in this chapter, we focus on evaluating the performance of the CDS method in terms of the accuracy of its approximation of the stochastic production planning problem. We do this using the standard problem setting assumed in the SCLSP literature with dynamic and uncorrelated demand. An application of the method to a setting with seasonal and correlated demand will be presented in Chapter 5. Specifically, we now investigate the performance of the CDS approach (in this section CDS in short) using the test instances presented in Helber et al. (2013) and compare it to that of the SCLSP-SCN approach with SRS and DS (in this section SRS and DS in short, respectively). A complete analysis should compare both the (expected) service levels achieved by the methods and their expected costs. However, comparing the expected costs of a method which underachieves the target for only one product with the expected costs of another method which misses the target for two products is not straightforward. If the first method leads to much higher costs, the better service-level performance might not be worth it. As a result, we choose to compare the methods by evaluating their performance based on the size of the expected over- or underachievement of the target service level for all products: the size of the overachievement will serve as an indicator of the expected costs.

In total, we consider $1,296$ test instances, each of which is characterized by a different combination of parameters, whose values are shown in Table 6. For every instance, we assume that the demand for each product-period combination is independent of other combinations and normally distributed with known mean $E[D_{kt}]$ and standard deviation σ_{kt}, which vary across instances. Further details on the parameters of the test instances can be found in Helber et al. (2013) and Appendix A.

Table 6: Parameters of the test instances – comparison of the SCN and CDS approximation of the SCLSP

Number of products	$K = \{5, 10, 20\}$
Number of time periods	$T = \{5, 10, 20\}$
Inter-period coefficient of variation of expected demand	$VC^{ip} = \{0.2, 0.3\}$
Demand coefficient of variation	$VC^{d} = \{0.1, 0.3\}$
Time between orders	$TBO = \{1, 2, 4\}$
Utilization due to processing	$Util = \{0.6, 0.75\}$
Setup time as fraction of period processing time	$ts^{rel} = \{0.0, 0.25\}$
Service-level target	$\delta = \{0.8, 0.9, 0.95\}$
Number of scenarios	$S = \{10, 30, 50\}$

Source: Helber et al. (2013)

We solve all problem instances using the SRS, DS and CDS approaches, each with 10, 30 and 50 scenarios, in order to compare their performances. We use different numbers of scenarios for each method because the quality of the solution and the solution time of all methods are strongly influenced by the number of scenarios used in the problem, which is always the case in scenario-based SP. Becauase of the mutual independence of demands, we note that the cumulative demand distributions needed by the CDS method can be easily obtained. For a given product k, the cumulative demand up to and including period t, CD_{kt}, is normally distributed, because it is the sum of normally distributed single-period demands, with mean $\mu_{CD_{kt}} = \sum_{\tau=1}^{t} E[D_{k\tau}]$ and standard deviation $\sigma_{CD_{kt}} = \sqrt{\sum_{\tau=1}^{t} \sigma_{kt}^2}$. Due to the size of some problem instances, we apply the F&O heuristic presented by Helber et al. (2013) to all instances to solve them in a reasonable time.

After solving all instances, we can evaluate the plans obtained by all the methods analytically, based on the true distribution of the random demand variables. Indeed, because all the random demand variables follow a known normal distribution, the expected backlog and expected inventory function values for a certain production quantity can be obtained exactly using the first-order loss function (see Helber et al., 2013). Therefore, running simulations to evaluate the performance of the methods is unnecessary.

We performed the analysis using an Intel Core CPU with 2.4 GHz and 16 GB of RAM. We coded and implemented the F&O heuristic in Python 3.7 and solved the subproblems using Gurobi 9.0. The optimization of each subproblem was stopped as soon as an optimality gap of 0.5% was reached or the time limit of 30 seconds was exceeded.

We first compare the different methods in terms of the average solution time of the test instances. In Table 7 we show, for all methods, the average run time of the F&O heuristic in CPU seconds as a function of the number of products, time periods and scenarios. As can be seen, there is no significant difference in solution time between the methods. In addition, as expected, the larger the number of periods, products and scenarios, the longer it takes to solve the problem for all methods.

As the starting point for the comparison of the performance of the methods, in Table 8 we show the average percentage of products over all instances for which the under and overachievement of the target service level fall in a given interval. For example, column "U≤0.01" shows, for a given method, the percentage of products over all instances which underachieve the target service level by at most one percentage point, whereas column "O≤0.01" shows the percentage of products over all instances that achieve the service level or overachieve it by at most one percentage point.

Figure 25 provides a visualization of the results of Table 8 with the use of relative frequency histograms. Specifically, each frequency plot is constructed by drawing, for each equal-length interval that represents the difference between the achieved and target service level, a rectangle with the interval as its base and the relative frequency of the

Table 7: Solution time (in CPU seconds)

			T=5	T=10	T=20
K=5	CDS	S=10	0.43	2.08	12.11
		S=30	1.28	6.49	43
		S=50	2.34	11.07	89.61
	DS	S=10	0.48	2.44	13.02
		S=30	1.45	7.27	43.98
		S=50	2.60	11.85	81.88
	SRS	S=10	0.47	2.48	12.83
		S=30	1.42	7.46	45.09
		S=50	2.56	12.05	86.73
K=10	CDS	S=10	1.42	7.79	59.18
		S=30	4.50	16.89	244.80
		S=50	8.09	28.36	550.35
	DS	S=10	1.57	8.80	68.05
		S=30	5.19	19.81	252.55
		S=50	9.24	28.38	503.61
	SRS	S=10	1.60	8.72	66.26
		S=30	5.00	18.87	253.51
		S=50	9.48	28.74	528.67
K=20	CDS	S=10	6.00	35.03	302.80
		S=30	16.54	64.05	1290.22
		S=50	29.34	137.07	3137.22
	DS	S=10	6.41	37.67	334.67
		S=30	19.82	66.86	1332.97
		S=50	31.36	133.55	2769.03
	SRS	S=10	6.66	37.94	343.48
		S=30	19.78	66.60	1370.92
		S=50	31.55	128.37	2858.50

interval over all products and instances as its height.

It is evident that as the number of scenarios increases, for both DS and SRS, the percentage of products that achieve a service level close to the target increases in the number of scenarios (the frequency "mass" of the intervals closer to 0% increases), indicating that the methods' solution quality improves. This is not a surprising result, as it is well known in SP that with sampling methods the quality of the solution improves as the number of scenarios increases, because a better representation of the underlying uncertainty is achieved. However, importantly, this improvement is not as significant for CDS, whose performance appears to be independent of the number of scenarios.

Moreover, the results show that DS and SRS clearly outperform CDS in terms of the percentage of products which meet the target service level (the sum of the last three columns of Table 8). However, CDS outperforms the other methods in terms of the percentage of products which miss the target by at most 1% (the sum of the last four

Table 8: Comparison of service level performance

		2%<U	1%<U≤2%	U≤1%	O≤1%	1%<O≤2%	2%<O
CDS	S=10	0.00%	0.00%	89.26%	3.43%	2.23%	5.07%
	S=30	0.00%	0.00%	90.34%	3.13%	2.23%	5.07%
	S=50	0.00%	0.00%	90.30%	3.23%	1.86%	4.62%
DS	S=10	0.13%	2.72%	60.44%	28.58%	2.72%	5.41%
	S=30	0.00%	0.34%	56.03%	36.56%	2.08%	4.99%
	S=50	0.00%	0.00%	53.33%	39.77%	2.06%	4.84%
SRS	S=10	6.98%	10.24%	33.79%	31.64%	8.85%	8.51%
	S=30	1.44%	6.33%	40.90%	39.26%	6.31%	5.77%
	S=50	0.37%	3.86%	43.16%	42.66%	4.87%	5.07%

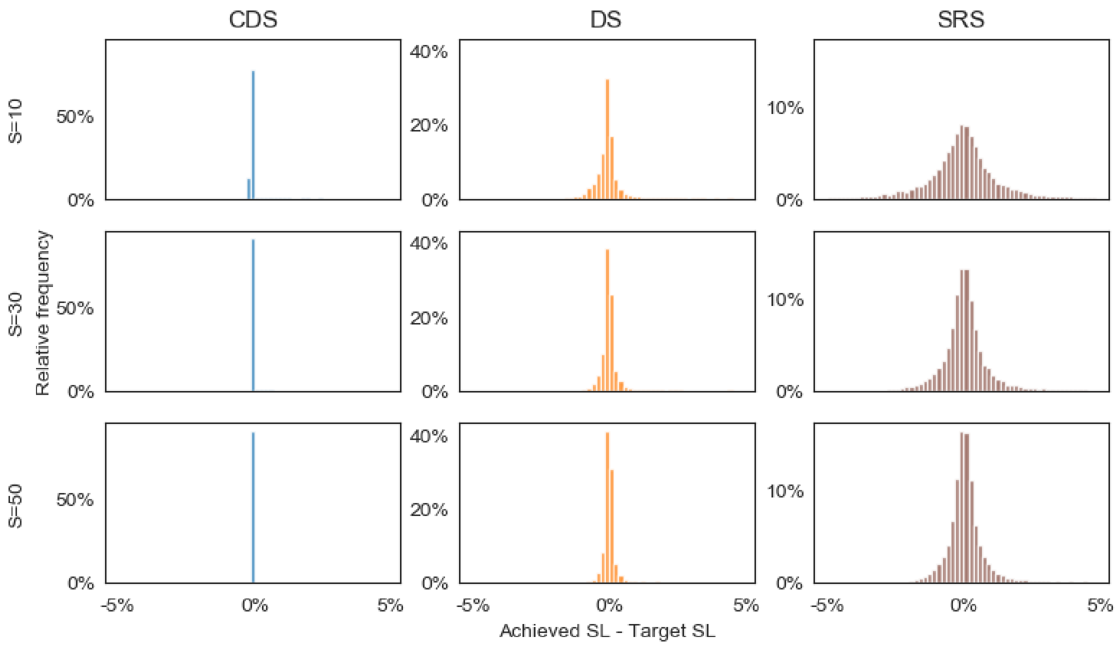

Figure 25: Comparison of service level performance

columns), even when only 10 scenarios are used to solve the problem. Therefore, if it is acceptable to miss the target by less than one percentage point, CDS provides a solution which is better than the one of SRS with 50 scenarios and as good as that of DS with 50 scenarios. Considering the increase in complexity of the model, which translates into a significant increase in solution time, as shown in Table 7, this is a remarkable result. Indeed, if we consider instances with 20 products and 20 time periods, we see that CDS with 10 scenarios can be solved approximately 9 times faster than both SRS and DS with 50 scenarios, clearly showing the advantage of CDS over the other methods in terms of computational effort. Concerning SRS, although the method outperforms CDS in terms of the percentage of products meeting the service-level target, the results show that it is outperformed by both CDS and DS in terms of the percentage of products which miss the target by more than 1%. The ability of the SRS method to limit the underachievement of the target improves along with the number of scenarios, as expected for completely random sampling. However, even with 50 scenarios, the method is less effective than the other two methods.

As an additional performance indicator, for all the methods and numbers of scenarios used, we calculate the average underachievement and overachievement of the target over all instances. We show the results using a bar chart in Figure 26, which illustrates that, for any given number of scenarios, CDS outperforms the other two methods both in terms of expected underachievement and expected overachievement. Moreover, the expected underachievement of CDS with 10 scenarios is only slightly larger than that of DS with 50 scenarios and much smaller than that of SRS with 50 scenarios. Finally, we see that the expected overachievement of CDS with 10 scenarios is below that of both DS and SRS with 50 scenarios. Overall, we can conclude that CDS outperforms the other two methods in terms of both the expected underachievement and overachievement.

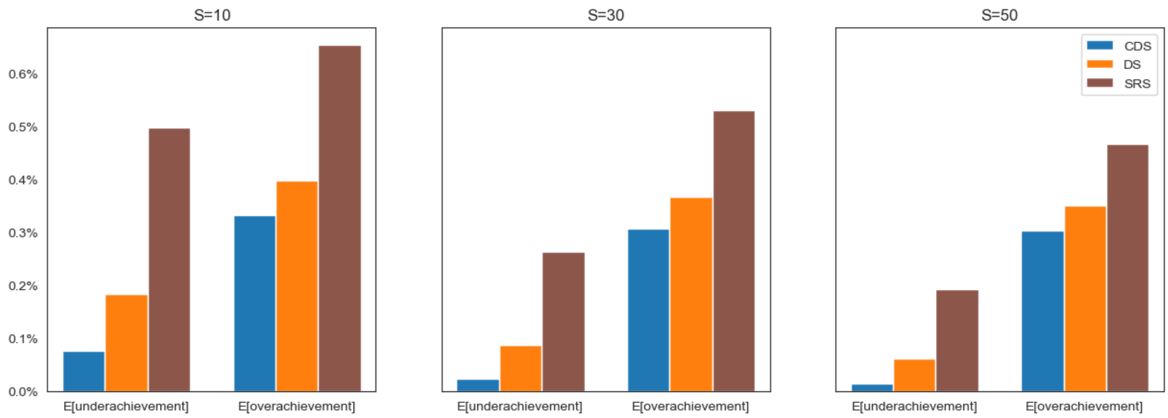

Figure 26: Comparison of average expected under and over achievement

An interesting characteristic of CDS is that it seems to systematically slightly underachieve the target, as can be seen from column "U≤0.01" of Table 8. As a reminder, given a production plan, the true δ service level of product k, denote it by δ_k^{true}, is defined as follows:

$$\delta_k^{true} = 1 - \frac{\sum_{t=1}^{T} E[BL_{kt}]}{\sum_{t=1}^{T} E[CD_{kt}]}. \tag{89}$$

In the SCLSP-CDS model, we approximate both the numerator and the denominator in the right-hand side using scenarios, obtaining the estimated service level

$$\delta_k = 1 - \frac{\sum_{t=1}^{T} \sum_{s=1}^{S} pr_s \cdot BL_{kts}}{\sum_{t=1}^{T} \sum_{s=1}^{S} pr_s \cdot CD_{kts}}. \tag{90}$$

In the case that demand for all products in all periods is normally distributed and that we use DS to obtain the scenarios in the CDS approach, it can easily be proven that the denominator of (90) estimates the denominator of (89) exactly. Therefore the systematic underestimation of the service level stems from a systematic underestimation of the expected backlog. In our numerical study, we have indeed observed that the expected backlog function for every product and period combination is always below the true one. As an example, in Figure 27 we show, for a normally distributed cumulative demand with mean 200 and standard deviation 50, the true expected backlog as a function of the cumulative production x and its approximation obtained with the CDS approach using $S = 10$. As can be seen, the CDS expected backlog function is always below the true one.

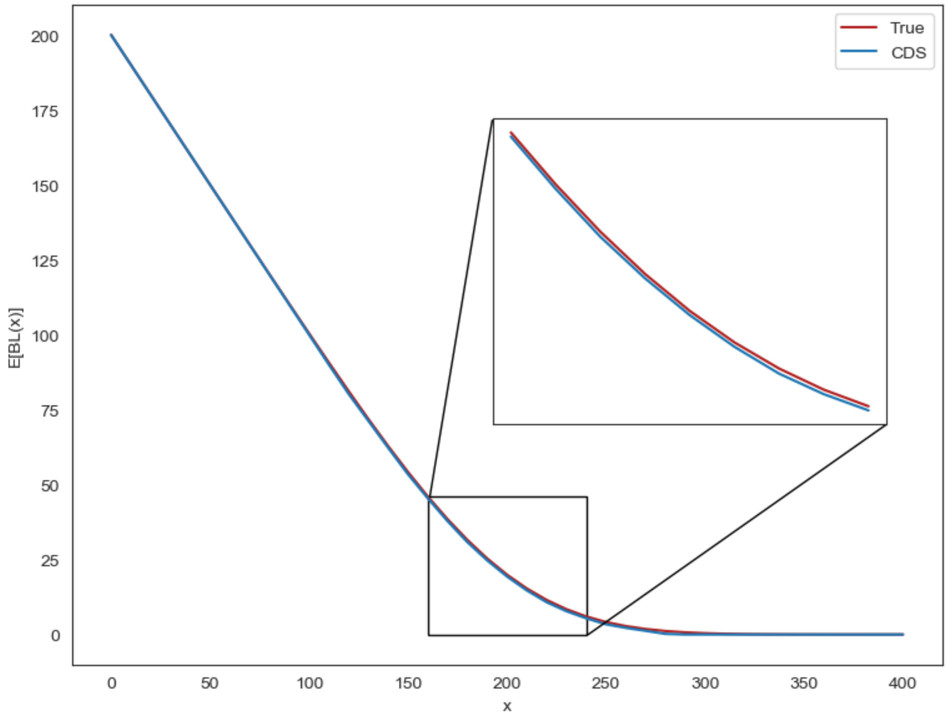

Figure 27: Comparison of true and CDS-approximated BL function

This apparent property suggests that we could obtain even better solutions for CDS by inflating the service-level target by a few percentage points. Table 9 shows the performance of all methods for all instances with $K = 5$, where CDS was solved increasing the original service-level targets by 0.3% to $\delta = \{0.803, 0.903, 0.953\}$. As can be seen, CDS now achieves the service-level target for 80% of the products with $S = 10$, and for 100% of the products with $S = 30$ and $S = 50$, thereby outperforming all methods. As Figure 28 shows, this comes at the price of a larger expected overachievement of the target and, therefore, expected costs, but this figure is still below that of SRS and only slightly above that of DS. It is important to note, however, that the "right" inflated target to use to reach the desired service level performance is problem-specific and difficult to determine a priori.

We now turn to analyze the effects of changing the parameters of the problem (listed in Table 6) on the quality of the solution for all the methods. To simplify the presentation of the results and the discussion, we focus on the service-level performance and use three performance indicators, following Helber et al. (2013): "SL", "SL 1%" and "SL 2%". These indicators represent the percentage of instances in which, in expectation, all

Table 9: Comparison of service level performance with inflated target for CDS

		2%<U	1%<U≤2%	U≤1%	O≤1%	1%<O≤2%	2%<O
CDS	S=10	0.00%	0.00%	19.26%	72.87%	2.64%	5.23%
	S=30	0.00%	0.00%	0.00%	92.92%	2.18%	4.91%
	S=50	0.00%	0.00%	0.00%	92.73%	1.99%	5.28%
DS	S=10	0.14%	3.10%	59.31%	28.24%	3.19%	6.02%
	S=30	0.00%	0.60%	56.67%	34.77%	2.64%	5.32%
	S=50	0.00%	0.00%	52.04%	39.91%	2.64%	5.42%
SRS	S=10	7.04%	9.95%	33.10%	32.59%	8.33%	8.98%
	S=30	1.39%	6.44%	39.72%	38.10%	7.59%	6.76%
	S=50	0.46%	3.19%	41.39%	43.33%	5.69%	5.93%

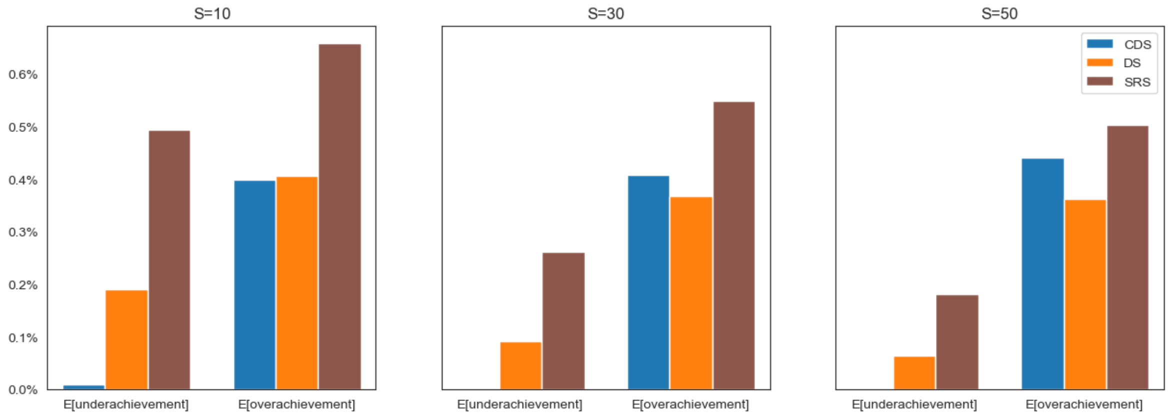

Figure 28: Comparison of average expected under and over achievement with inflated target for CDS

products meet the service-level target, underachieve it by at most 1% and 2%, respectively. Table 10 shows the effect of changing the coefficient of variation of demand, VC^d. From the decision maker's perspective, this parameter can be thought of as the degree of forecast accuracy, i.e. the larger the coefficient of variation of demand, the more inaccurate the results are. The results show that SL slightly decreases for all methods as the coefficient of variation of demand increases. In addition, SL 1% decreases for SRS and DS, however, this decrease is less important for larger scenario sets (only for DS with $S = 50$ the SL 1% performance remains unchanged). Instead, for CDS, the SL 1% performance does not decrease in the coefficient of variation of demand, independent of the number of scenarios used by the method. This clearly shows that, from a managerial perspective, in terms of the SL1 % indicator, CDS is robust against forecast inaccuracy and computationally efficient, because decreasing the number of scenarios does not affect the method's performance.

Table 11 shows the effects of varying the service-level target. Focusing again on the SL 1% performance indicator, for any number of scenarios, SRS performs consistently poorly, whereas the performance of DS with 10 and 30 scenarios clearly worsens as the target increases (for S=50 the performance remains equal). Instead, the performance of

Table 10: Comparison of service level performance for different values of the coefficient of variation of demand

			SL	SL 1%	SL 2%
CDS	S=10	$VC^d = 0.1$	0.93%	100.00%	100.00%
		$VC^d = 0.3$	0.31%	100.00%	100.00%
	S=30	$VC^d = 0.1$	0.62%	100.00%	100.00%
		$VC^d = 0.3$	0.15%	100.00%	100.00%
	S=50	$VC^d = 0.1$	0.15%	100.00%	100.00%
		$VC^d = 0.3$	0.15%	100.00%	100.00%
DS	S=10	$VC^d = 0.1$	2.47%	100.00%	100.00%
		$VC^d = 0.3$	0.93%	68.83%	97.53%
	S=30	$VC^d = 0.1$	2.31%	100.00%	100.00%
		$VC^d = 0.3$	0.77%	95.06%	100.00%
	S=50	$VC^d = 0.1$	2.31%	100.00%	100.00%
		$VC^d = 0.3$	1.39%	100.00%	100.00%
SRS	S=10	$VC^d = 0.1$	3.09%	62.81%	96.14%
		$VC^d = 0.3$	0.77%	7.87%	29.01%
	S=30	$VC^d = 0.1$	2.62%	92.28%	99.85%
		$VC^d = 0.3$	1.08%	27.31%	74.85%
	S=50	$VC^d = 0.1$	3.86%	98.15%	100.00%
		$VC^d = 0.3$	1.70%	44.75%	92.75%

CDS is perfect independent of the target and the number of scenarios. This shows that CDS has the desired property of being robust to changes in the target service level, as opposed to DS. In addition, when compared to DS and SRS, it is evident that CDS can again provide high quality solutions even with only 10 scenarios.

Finally, Table 12 shows the effect of changing the TBO parameter. As a reminder, this parameter influences the setup costs, with a larger TBO indicating higher setup costs. Continuing to focus on the SL 1% performance of the methods, we can observe that the performance of both DS and SRS improves in the setup costs, whereas that of CDS stays perfect, again illustrating the robustness of the CDS method to the problem's parameters. We found that the rest of the problem's parameters did not significantly influence the performance of the different methods.

Table 11: Comparison of service level performance for different values of the target service level

			SL	SL 1%	SL 2%
CDS	S=10	$\delta = 0.8$	1.85%	100.00%	100.00%
		$\delta = 0.9$	0.00%	100.00%	100.00%
		$\delta = 0.95$	0.00%	100.00%	100.00%
	S=30	$\delta = 0.8$	1.16%	100.00%	100.00%
		$\delta = 0.9$	0.00%	100.00%	100.00%
		$\delta = 0.95$	0.00%	100.00%	100.00%
	S=50	$\delta = 0.8$	1.16%	100.00%	100.00%
		$\delta = 0.9$	0.00%	100.00%	100.00%
		$\delta = 0.95$	0.00%	100.00%	100.00%
DS	S=10	$\delta = 0.8$	4.40%	93.06%	99.77%
		$\delta = 0.9$	0.69%	82.64%	98.61%
		$\delta = 0.95$	0.00%	77.55%	97.92%
	S=30	$\delta = 0.8$	4.17%	99.77%	100.00%
		$\delta = 0.9$	0.23%	97.45%	100.00%
		$\delta = 0.95$	0.23%	95.37%	100.00%
	S=50	$\delta = 0.8$	4.63%	100.00%	100.00%
		$\delta = 0.9$	0.46%	100.00%	100.00%
		$\delta = 0.95$	0.46%	100.00%	100.00%
SRS	S=10	$\delta = 0.8$	4.63%	37.96%	62.96%
		$\delta = 0.9$	0.93%	33.10%	61.11%
		$\delta = 0.95$	0.23%	34.95%	63.66%
	S=30	$\delta = 0.8$	4.17%	59.72%	85.88%
		$\delta = 0.9$	0.46%	59.26%	85.88%
		$\delta = 0.95$	0.93%	60.42%	90.28%
	S=50	$\delta = 0.8$	5.56%	69.68%	95.60%
		$\delta = 0.9$	1.39%	69.91%	95.60%
		$\delta = 0.95$	1.39%	74.77%	97.92%

Table 12: Comparison of service level performance for different values of the time between orders

			SL	SL 1%	SL 2%
CDS	S=10	TBO= 1	1.62%	100.00%	100.00%
		TBO= 2	0.23%	100.00%	100.00%
		TBO= 4	0.00%	100.00%	100.00%
	S=30	TBO= 1	1.16%	100.00%	100.00%
		TBO= 2	0.00%	100.00%	100.00%
		TBO= 4	0.00%	100.00%	100.00%
	S=50	TBO= 1	1.16%	100.00%	100.00%
		TBO= 2	0.00%	100.00%	100.00%
		TBO= 4	0.00%	100.00%	100.00%
DS	S=10	TBO= 1	3.01%	75.46%	96.53%
		TBO= 2	1.16%	84.26%	99.77%
		TBO= 4	0.93%	93.52%	100.00%
	S=30	TBO= 1	2.78%	94.21%	100.00%
		TBO= 2	0.46%	98.84%	100.00%
		TBO= 4	1.39%	99.54%	100.00%
	S=50	TBO= 1	3.01%	100.00%	100.00%
		TBO= 2	0.46%	100.00%	100.00%
		TBO= 4	2.08%	100.00%	100.00%
SRS	S=10	TBO= 1	2.55%	26.39%	52.08%
		TBO= 2	1.62%	34.49%	65.05%
		TBO= 4	1.62%	45.14%	70.60%
	S=30	TBO= 1	3.70%	52.08%	80.79%
		TBO= 2	0.23%	58.10%	86.11%
		TBO= 4	1.62%	69.21%	95.14%
	S=50	TBO= 1	3.70%	63.66%	92.82%
		TBO= 2	2.08%	69.91%	97.22%
		TBO= 4	2.55%	80.79%	99.07%

To summarize, the numerical study presented in this section shows that the newly developed CDS approach outperforms the standard scenario-approximation methods proposed in the literature in many aspects, even in the absence of demand dependencies. Indeed, the average percentage of products that miss the target service level by at most 1% is largest for CDS, and both the expected underachievement and overachievement of the target is lowest for CDS. Moreover, the study highlighted the computational efficiency of CDS compared to other approaches, because in terms of solution quality the results of the method with 10 scenarios are very close to the results when three or five times the number of scenarios are used. Finally, we showed that, as opposed to SRS and DS, the new method, even with only 10 scenarios, is robust to changes in all parameters of the problem.

4.6 Summary

In this chapter, we presented a new formulation of the scenario approximation of the SCLSP under a static uncertainty strategy using δ service-level constraints. The core difference between the newly proposed SCLSP-CDS model and the standard scenario approximation of the problem, SCLSP-SCN, is that the former uses scenarios of cumulative demand for each product-period combination to estimate expected backlogs and inventory, as opposed to paths of demand in a scenario tree/fan as the latter. In complex application settings with demand dependencies, this difference in the nature of scenarios significantly simplifies the scenario-generation process while leading to a correct estimation of expected inventory and backlog. Indeed, using the SCLSP-CDS scenario approximation, first, it is unnecessary to estimate the correct correlations between the demand of different (k, t) pairs in the pre-scenario-generation phase. Second, generating the scenarios for the cumulative model consists of sampling from marginal (cumulative) demand distributions, instead of a joint distribution with dependencies between the random variables. Moreover, we emphasize that estimating the cumulative demand distribution for a given product-period combination requires the same amount of effort as estimating the period-specific demand distribution for the same product-period combination, which is a necessary step in other standard scenario-generation procedures.

In our numerical study we considered the case of independent demands and showed that even in this setting the newly developed approach outperforms other common scenario-approximation methods to solve the SCLSP in terms of most performance measures. In addition, it is important to note that the new method is computationally efficient, i.e. it is able to achieve good quality solutions with a much smaller number of scenarios than the other methods. Indeed, in the analyzed problem instances, the CDS approach provides an excellent approximation of the expected backlog and inventory functions, even when a limited number of scenarios is used. Finally, we have also shown the noteworthy property of the CDS method of being robust to changes in the parameters of the problem, namely the coefficient of variation of demand, the service-level target and the time-between-orders. We defer an in-depth investigation into the reasons for the computational efficiency and the robustness of the CDS approach to future studies.

We think that a comparison in a setting with dependent demands between the SCLSP-CDS approach and the standard SCLSP-SCN approach, with the latter using other sampling or statistical scenario-generation methods able to deal with demand dependencies, would produce further valuable insights. Therefore, we propose this task as a future research opportunity. In the next chapter of this thesis, we move in this direction by presenting an application of the CDS approximation to a different production setting with demand dependencies. Another opportunity of interest for future research is to evaluate the performance of the SCLSP-CDS approach in a RH planning strategy, which is a frequently used method to deal with demand uncertainty in practice.

To conclude, notwithstanding their limitations, we believe that the tools presented in this chapter have the potential of being highly valuable for inventory planning in many practical settings with demand dependencies.

5 Inventory analysis for an agrochemical company: a real-world case study

5.1 Introduction and research objectives

In this chapter, we use the cumulative demand scenario (CDS) approach introduced in Chapter 4 to study the mid-term inventory problem for crop protection products (CPPs) faced by an agrochemical company. For reasons of confidentiality, we cannot disclose the name of the company and will instead refer to it as Agro Chem (AC). The investigated setting is not specific to the industry partner though, but it is typical for a large size agrochemical producer. Specifically, in this study, we consider the inventory problem for the entire portfolio of CPPs produced at a formulation plant of AC. The demand for these products, as is typical for agrochemical products, is seasonal and highly uncertain. Indeed, the demand for each of these CPPs occurs only in a limited time period during the year and is highly uncertain, both in terms of timing and quantity, due to the many factors influencing the buying decisions, most notably the weather. Accordingly, there are strong dependencies over time and across products that affect demand, and these dependencies must be considered when making inventory decisions. Moreover, the production processes of the considered CPPs are highly complex due to the nature and use of the products. This complexity makes it necessary to make lot-sizing and scheduling decisions simultaneously in order to derive feasible and cost-efficient plans. Thus, in contrast to Chapter 4, we now analyze a more complex real-world setting. In particular, in this chapter we develop a stochastic lot-sizing and scheduling model, which we solve using the CDS approach, and apply it to a setting with seasonal, uncertain and dependent demand.

Currently, AC solves its mid-term inventory planning problem in two steps, which are performed sequentially. In the first step, expected future demands are inflated by experts to consider demand uncertainty and AC's target service level. These decisions essentially determine safety stock quantities and determine the inventory targets for all products in all periods of the planning horizon considered. In the second step, a deterministic planning problem, which aims to meet these inventory targets, is solved. This type of sequential planning approach is very common in practice and in the literature. However, it is known to lead to suboptimal results, because the interaction between safety stock and cycle stock is disregarded. Moreover, in the specific case of AC, the first step heavily relies on the experts' subjective judgments. The company recognizes the importance of demand uncertainty and requests a tool able to provide a clear picture of how this uncertainty influences its inventory costs and service-level performance. Moreover, AC is interested in evaluating their currently used planning approach to determine its ability to manage demand uncertainty.

Therefore, in this study we have two objectives:

1. Provide the company with insights into the impact of demand uncertainty on their mid-term inventory/production planning problem;

2. Assess the effectiveness of the methods currently employed by the company in their decision-making processes to deal with demand uncertainty.

To reach these two closely connected goals, we use a stochastic general lotsizing and scheduling problem (SGLSP) to model the inventory problem faced by the company, and use the CDS approach to approximate and solve this problem, because of the uncertain seasonal demand of the products considered. This integrated model, as opposed to the sequential procedure currently implemented by the company, uses the available statistical information on the demand process to make simultaneous lot-sizing, scheduling and safety stock decisions.

Therefore, in this chapter we contribute to the literature on inventory planning in settings with demand seasonality and uncertainty, particularly the agrochemical industry that was presented in Chapter 2, by presenting a real-world case study conducted at AC.

The rest of this chapter is organized as follows. In Section 5.2 we present the problem setting in detail, as well as the stochastic model used to conduct the study. In Section 5.3 we explain how we derive the input data of the model, with a particular focus on the probability distribution of demand. The results are presented in Section 5.4. Finally, Section 5.5 concludes the chapter with a summary of the findings and future research opportunities.

5.2 Problem setting and model formulation

In this section, we present the setting of the analyzed inventory planning problem and the mathematical optimization model developed to model it. In Section 5.4 we then use this model to conduct an analysis to achieve the two objectives of this chapter. To protect the company's interests, we cannot fully disclose all the details pertaining to the problem setting; nonetheless, the details we present provide a clear description of the analyzed problem.

Our focus lies on the mid-term inventory management problem that a single plant of AC experiences. This consists of developing a rough-cut production and inventory plan which enables the plant to meet the company's goals for a given year under demand uncertainty. The plant considered is a formulation plant for a relatively wide range of CPPs, sold mostly to customers located in the same geographical region as the plant. The production process to manufacture these CPPs follows the steps previously described in Chapter 2. First the active ingredients (AIs) are synthesized in dedicated plants, and then formulation takes place in the plant under consideration. After the formulation process, CPPs are bottled and labeled according to the market in which the products are finally sold (we name this final country-specific form in which the CPPs are sold "articles"). As a reminder, the typical supply chain (SC) of CPPs is shown in Figure 29.

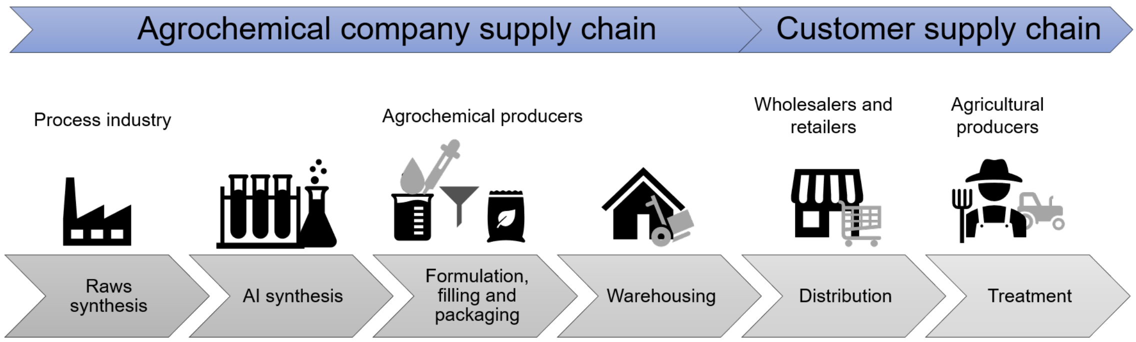

Figure 29: Typical supply chain of crop protection products. Adapted from Fritz and Hausen (2009)

The plant consists of less than 10 production resources, each of which can produce a different (potentially overlapping) set of the approximately 40 CPPs formulated in the plant. These CPPs are sold in the form of country-specific articles, which in total are in the hundreds. These products have a seasonal and uncertain demand, but every product has a potentially different season, because the demand depends on the agricultural cycles of crops and the life cycles of the pests that the product is designed to eliminate. However, there is still an overall seasonal demand pattern, i.e. an aggregate off-season and peak season for all products. We consider the firm's planning problem at the beginning of the first month of this "macro off-season", when it needs to decide how to build up stock for the upcoming macro peak season. Uncertainty exists in both demand timing and scale, and it is therefore characterized by demand dependencies, both between periods and products.

An important feature of the formulation process of the CPPs considered which significantly complicates production planning is the sequence-dependence of the setup operations, i.e. the fact that the duration and/or cost of setting up a resource to produce a given product depends on the product that was produced before. This is due to the nature of the products: the need to avoid cross-contamination between different CPPs, mostly those using different AIs, requires the resources to be thoroughly cleaned when a changeover is performed, i.e. when two different products are produced sequentially on the same resource.

First we conceptually analyze the choices made to model the plant's inventory problem that must be solved by the planners, and then present the mathematical formulation of the model.

We focus on and model the formulation step of the production process. In addition, we model the inventory process and consider the transportation of the products to the destination warehouses implicitly through appropriate lead times. However, due to the mid-term nature of the problem, for simplification purposes we model the production, inventory and demand processes at the level of CPPs, instead of the country-specific articles. The production of AIs, which serve as input in the formulation plant, is not considered in this study. In particular, we assume that there is a sufficient supply of AIs to sustain any (realistic) production plan. We also disregard bottling and labeling operations, although these could be considered in a future study.

While on the one hand we are looking at a mid-term planning problem, and therefore, we are not interested in modeling the dynamics of production, inventory and demand in detail, on the other hand we have to recognize that sequence-dependent setups require a more complex model than the stochastic capacitated lot-sizing problem (SCLSP) analyzed in Chapter 4. Usually, setups and sequence-dependent setups are not modeled in mid-term planning problems, but in certain settings, such as the process industry analyzed here, their impact is significant enough to justify the additional model complexity caused by their inclusion (Albrecht et al., 2015). We, therefore, include them in the production model presented in this chapter because a wrong production sequence can result in highly significant changeover times (in the order of days). Because the proposed model could be used, after appropriate modifications, to support future planning efforts, including this very challenging feature into the problem would also facilitate the acceptance of the model by the planning experts, as well as produce more realistic plans which can then be used as input for more detailed short-term planning.

We choose the general lot-sizing and scheduling problem (GLSP) (Fleischmann and Meyr, 1997) to model this planning problem, because of its ability to elegantly model sequence-dependent setups and its flexibility in accommodating additional characteristics of the problem. However, in view of the uncertainty and seasonality of demand for the products under consideration, an appropriate stochastic version of the GLSP is necessary. Therefore, we develop a stochastic GLSP and solve it by using the cumulative demand scenario approach presented in Chapter 4.

As mentioned, we consider a planning horizon of one year, and use monthly time buckets. We choose the length of one year because of the need to consider at least an entire season in the problem, while at the same time limiting the increase in computational effort caused by considering longer time horizons. However, "truncating" the time horizon might have negative consequences, as the optimal decisions within the horizon might be dependent on the demand in periods beyond it. In particular, these consequences, known as end-of-horizon (EOH) effects, are that the optimal solution consists of not performing any setups in the last periods of the planning horizon and ending with zero inventories, because these inventories do not have any purpose beyond the horizon (Lang, 2010). In our case, inventory at the end of the planning horizon will be needed to satisfy the demand in the first periods of the next off-season. To avoid EOH effects, several techniques have been proposed in the literature, e.g. setting final inventory targets or decreasing the setup costs of lots in the final periods of the planning horizon (Lang, 2010). We opt to implement ending inventory constraints, namely forcing the expected ending net inventory to be at least larger than the starting inventory. This creates a cyclical plan which ensures that, in expectation, we can meet the service level also in the next season. Moreover, as opposed to production planning models, we define initial inventory as a decision variable in our model. The purpose of this is to ensure that the analysis for a given season is independent of the starting conditions of the inventory system. Indeed, when comparing the performance of the developed stochastic model with the currently used inventory management tools (one of our two research objectives), the "wrong" starting inventory might obscure some of the advantages that one method has over the other.

In addition, we also consider the lead time needed for the products to be transported in their intended market by requiring products to be produced by the end of the period preceding the one in which they are sold. We complete the cyclical schedule mentioned above by making the production quantity of the last period of the planning horizon, together with initial inventory, available to serve the demand in the first period.

As discussed in Chapter 4, there are different strategies to deal with demand uncertainty. In our model we use the static uncertainty strategy. This means that all production decisions, i.e. the production schedule and lot sizes, are fixed at the start of the planning horizon.

The objective of the model is to minimize (expected) inventory costs subject to meeting a target service level. Furthermore, we model sequence-dependent setups only in terms of time, and not in terms of costs in the objective function. However, we note that, apart from the operational costs for performing the changeover operation, a major component of setup costs is the opportunity cost of the setup. Indeed, the time used for performing the changeover could be used to produce the same or other products closer to the selling season, thus saving inventory costs. Therefore, by considering sequence-dependent setup times, this important element of setup costs is taken into account and traded off with inventory costs (the same approach is used in Tavaghof-Gigloo, 2019). In the objective function, we also do not consider other cost components, e.g. production, overtime or transportation costs, because these are beyond the scope of this study.

We assume that unmet demand is backordered and we limit the amount of backorders by enforcing a service-level target. As presented in Chapter 4, there are many different service level metrics used for this purpose. We use the already introduced δ service level metric, because of its properties and the ease of its implementation in the CDS approach, which will be used to solve the analyzed stochastic planning problem. In the literature, the service level is usually measured for each product individually, which potentially enables setting a different target per product. Alternatively, companies might set targets at the product group level, or even a single overall target, depending on their strategic goals. In the initial formulation of the stochastic model, we measure the service level over all products, as we find that this simplifies the discussion of the results with our research partner. However, we modify this in a subsequent extension of the model to provide further insights. If there is a single “aggregate” service-level target, then a poor service-level performance for one product can be compensated for by a good service-level performance of another product.

In accordance with the problem setting and the modeling choices explained in the previous paragraphs, we finally present the mathematical formulation of the SGLSP which we use for the analysis. The GLSP is a hybrid lot-sizing model, because it uses both macro-periods, like big bucket models, and micro-periods, like small bucket models. Formally, the planning horizon, equal to one year, is divided into T macro-periods, each equal to one month, and each macro-period t is divided into a fixed number of micro-periods (to simplify the notation of the model, let $\mathcal{N}$ denote the set of all chronologically ordered micro-periods of the planning horizon, and $\mathcal{N}_t = \{N_{t-1}+1, \ldots, N_t\}$ the subset of micro-periods in macro-period t, where N_t is the last micro-period in t). Similar to most small bucket models, it is assumed that at most one product can be produced per micro-period; given that the number of micro-periods in a given macro-period is fixed, the total maximum number of products that can be produced, and therefore the number of changeovers that can be performed, in a macro-period is also fixed and equal to the number of micro-periods. However, although the number of micro-periods in a given macro-period is fixed in advance by the planner, their lengths are not fixed, but are implicit decision variables of the model: the length of a micro-period is equal to the time spent processing the product produced in the considered micro-period (naturally there can be idle micro-periods in which nothing is produced, and the assumption is that the setup state is conserved in this case). As explained by Fleischmann and Meyr (1997), the “external dynamics” of the production/inventory system are modeled in discrete time (the macro-periods), i.e. the demand and the holding costs are specified and occur at discrete points in time, whereas the “internal dynamics” are modeled in continuous time (the micro-periods), i.e. the changeovers in production can happen at any time within the planning horizon.

Uncertainty of demand is modeled using scenarios following the cumulative view approach presented in Chapter 4. Formally, using the notation in Table 13, the CDS approximation of the SGLSP is defined as follows:

SGLSP-CDS Model

$$\text{Min} \sum_{k\in\mathcal{K}} \sum_{t\in\mathcal{T}} h_k \cdot \Big(\sum_{s\in\mathcal{S}} pr_s \cdot IP_{kts} \Big) + \sum_{k\in\mathcal{K}} \sum_{n\in\mathcal{N}} \sum_{m\in\mathcal{M}} h_k \cdot x_{knm} \tag{91}$$

Subject to

$$\sum_{n\in\mathcal{N}_t}\sum_{k\in\mathcal{K}} pt_{km}\cdot x_{knm} + \sum_{n\in\mathcal{N}_t}\sum_{k\in\mathcal{K}}\sum_{j\in\mathcal{K}} st_{kjm}\cdot z_{kjnm} \le b_{mt} \qquad \forall m\in\mathcal{M},\ \forall t\in\mathcal{T} \tag{92}$$

$$x_{knm} \le \frac{b_{mt}}{pt_{km}}\cdot y_{knm} \qquad \forall k\in\mathcal{K},\ \forall m\in\mathcal{M},\ \forall n\in\mathcal{N} \tag{93}$$

$$x_{knm} \ge l_{km}\cdot(y_{knm} - y_{k,n-1,m}) \qquad \forall k\in\mathcal{K},\ \forall m\in\mathcal{M},\ \forall n\in\mathcal{N} \tag{94}$$

$$\sum_{k\in\mathcal{K}} y_{knm} = 1 \qquad \forall m\in\mathcal{M},\ \forall n\in\mathcal{N} \tag{95}$$

$$z_{kjnm} \ge y_{k,n-1,m} + y_{jnm} - 1 \qquad \forall k,j\in\mathcal{K},\ \forall m\in\mathcal{M},\ \forall n\in\mathcal{N} \tag{96}$$

$$NI_{kts} = IP_{k0} + \sum_{n\in\mathcal{N}_T}\sum_{m\in\mathcal{M}} x_{knm} - CD_{kts} \qquad \forall k\in\mathcal{K},\ t=1,\ \forall s\in\mathcal{S} \tag{97}$$

$$NI_{kts} = IP_{k0} + \sum_{n\in\mathcal{N}_T}\sum_{m\in\mathcal{M}} x_{knm} + \sum_{\tau=1}^{t-1}\sum_{n\in N_\tau}\sum_{m\in\mathcal{M}} x_{knm} - CD_{kts} \qquad \forall k\in\mathcal{K},\ t\ge 2,\ \forall s\in\mathcal{S} \tag{98}$$

$$IP_{kts} \ge NI_{kts} \qquad \forall k\in\mathcal{K},\ \forall t\in\mathcal{T},\ \forall s\in\mathcal{S} \tag{99}$$

$$BL_{kts} \ge -NI_{kts} \qquad \forall k\in\mathcal{K},\ \forall t\in\mathcal{T},\ \forall s\in\mathcal{S} \tag{100}$$

$$\sum_{k\in\mathcal{K}}\sum_{t\in\mathcal{T}}\sum_{s\in\mathcal{S}} pr_s\cdot BL_{kts} \le (1-\delta)\cdot\sum_{k\in\mathcal{K}}\sum_{t\in\mathcal{T}}\big(E[CD_{kt}]\big) \tag{101}$$

$$\sum_{n\in\mathcal{N}}\sum_{m\in\mathcal{M}} x_{knm} \ge E[CD_{kT}] \qquad \forall k\in\mathcal{K} \tag{102}$$

$$x_{knm} \ge 0 \qquad \forall k\in\mathcal{K},\ \forall m\in\mathcal{M},\ \forall n\in\mathcal{N} \tag{103}$$

$$y_{knm} \in \{0,1\} \qquad \forall k\in\mathcal{K},\ \forall m\in\mathcal{M},\ \forall n\in\mathcal{N} \tag{104}$$

$$z_{kjnm} \ge 0 \qquad \forall k,j\in\mathcal{K},\ \forall m\in\mathcal{M},\ \forall n\in\mathcal{N} \tag{105}$$

$$IP_{kts}, BL_{kts} \geq 0 \qquad \forall k \in \mathcal{K},\ \forall t \in \mathcal{T},\ \forall s \in \mathcal{S}. \tag{106}$$

The objective (91) is to minimize the expected costs of holding inventory, including pipeline stock (the second term of the objective function). Constraints (92)-(96) are the constraints defining the production process in the standard (deterministic) GLSP. Constraints (92) ensure that the production time plus the changeover times performed in any machine in every period does not exceed the capacity of the machine in that period. Constraints (93) ensure that the production of a product in a given machine can only occur if the machine is set up for the respective product. Constraints (94) make sure that every time there is a changeover in a machine to a different product, its minimum batch size is produced. This constraint is necessary not only for technical reasons pertaining to the production process, but also to avoid the possibility of performing a changeover to a product while producing no units of it just to take advantage of the non-triangular nature of the setup matrix. Constraints (95) ensure that each machine can be set up for a single product per micro-period. Constraints (96) define the changeover variables. Constraints (97)-(100) define, for each scenario, the net inventory, the physical inventory and the backlog of a product at the end of each period, respectively. These are the stochastic counterpart of the inventory balance constraints in the standard deterministic GLSP under the "cumulative demand view" of the CDS approach developed in Chapter 4. Constraint (101) is the aggregate δ service-level constraint. Constraints (102) ensure that the cumulative production over the entire planning horizon is larger than or equal to the cumulative demand. These constraints not only serve the same purpose as in the SCLSP-CDS model introduced in Chapter 4, as we again use the δ service-level measure, but also ensure that the expected ending net inventory is larger than or equal to the starting inventory, i.e. they are the EOH constraints. Constraints (103)-(106) define the domain of the decision variables of the model.

This model integrates the lot-sizing, scheduling and safety stock problems, thus considering the interaction of these decisions.

5.3 Derivation of the input data

We conduct our analysis under real-world conditions by evaluating the performance of the SGLSP-CDS model, and other models introduced later in this chapter to reach our research objectives, in three past seasons, named "Season 1", "Season 2" and "Season 3". Specifically, for each of the three seasons, we assume that we are in the first month of the macro off-season and have to make the inventory planning decision for the following year under uncertainty. After solving the models, we can then derive both their expected performance, i.e. according to the assumed probability distribution of demand, and their

Table 13: SGLSP model – notation

	Sets
$\mathcal{K} = \{1, \dots, K\}$	Set of products
$\mathcal{M} = \{1, \dots, M\}$	Set of machines
$\mathcal{T} = \{1, \dots, T\}$	Set of macro-periods
$\mathcal{N}_t = \{N_{t-1}+1, \dots, N_t\}$	Set of micro-periods in macro-period $t \in \mathcal{T}$
$\mathcal{N} = \{1, \dots, N\}$	Chronologically ordered set of all micro-periods
$\mathcal{S} = \{1, \dots, S\}$	Set of scenarios
	Parameters
b_{mt}	Production capacity of machine $m \in \mathcal{M}$ in macro-period $t \in \mathcal{T}$
pt_{km}	Unit processing time of product $k \in \mathcal{K}$ in machine $m \in \mathcal{M}$
st_{kjm}	Setup time of a changeover from product $k \in \mathcal{K}$ to product $j \in \mathcal{K}$ in machine $m \in \mathcal{M}$
h_k	Unit holding costs per macro-period for product $k \in \mathcal{K}$
l_{km}	minimum lot size of product $k \in K$ (units) in machine $m \in \mathcal{M}$
pr_s	probability of occurrence of scenario $s \in \mathcal{S}$
CD_{kts}	Cumulative demand of product $k \in \mathcal{K}$ up to macro-period $t \in \mathcal{T}$ in scenario $s \in \mathcal{S}$
$E[CD_{kt}]$	Expected cumulative demand of product $k \in \mathcal{K}$ up to macro-period $t \in \mathcal{T}$
δ	Delta service-level target
	Scenario-dependent decision variables
NI_{kts}	Net inventory of product $k \in \mathcal{K}$ at the end of macro-period $t \in \mathcal{T}$ in scenario $s \in \mathcal{S}$
IP_{kts}	Physical inventory of product $k \in \mathcal{K}$ at the end of macro-period $t \in \mathcal{T}$ in scenario $s \in \mathcal{S}$
BL_{kts}	Backlog of product $k \in \mathcal{K}$ at the end of macro-period $t \in \mathcal{T}$ in scenario $s \in \mathcal{S}$
	Scenario-independent decision variables
x_{knm}	Lot size of product $k \in \mathcal{K}$ in micro-period $n \in \mathcal{N}$ in machine $m \in \mathcal{M}$
IP_{k0}	Initial physical inventory of product $k \in \mathcal{K}$
y_{knm}	Binary setup variable of product $k \in \mathcal{K}$ in micro-period $n \in \mathcal{N}$ in machine $m \in \mathcal{M}$
z_{kjnm}	changeover variable equal to 1 if there is a changeover from product $k \in \mathcal{K}$ to $j \in \mathcal{K}$ in micro-period $n \in \mathcal{N}$ in machine $m \in \mathcal{M}$; 0 otherwise

actual performance, i.e. according to the actual demand realization in the season. In this section, we explain how we derive the input parameters of the SGLSP-CDS model, with a particular focus on the scenarios of cumulative demand.

Several parameters of the problem are relatively easy to obtain, such as the production capacity of the machines, the unit processing times of products, the changeover times, the unit holding costs and the minimum products' lot sizes. Defining a single service-level target for a stochastic planning problem is not easy from a practical point of view, so we vary this number and solve multiple optimization problems, which also allows us to highlight the trade-off between inventory costs and service level. The most challenging input parameters to obtain are the scenarios of cumulative demand for each product-period combination. To this end, we first need to estimate the cumulative demand distribution from which we can then sample to obtain the scenarios. The data at our disposal to reach this goal are historical sales and forecasts for a limited number of past seasons.

As a first step towards estimating the demand distribution, we assume that past sales values are a good representation of past actual demand, because the latter is unavailable. The classical approach used in the literature is to use past sales/demand data to derive a demand distribution for the future. However, this approach is problematic in this setting, because of the limited number of historical data points available and the nature of the market for CPPs. To recall, given the yearly seasonality of demand, even if 10 years of data were available, the cumulative demand distribution for each product-period combination would still need to be derived from only 10 data points, which are insufficient to obtain a well-defined demand distribution. Moreover, in our setting, data points taken from far back in time most likely are not representative of the demand in future seasons, due to a different portfolio of products, different customers, different competitors' portfolios, etc.

To solve the representativeness problem, we note that, apart from past sales data, as mentioned, in this particular setting we also have access to past forecasts data: combined with demand data, this provides us with a time series of forecast errors, which gives us information on demand uncertainty. We, therefore, derive the demand distribution for the planning problem of each analyzed season starting from the historical forecasts for the season and the historical information on forecast errors. Accordingly, we make the seemingly reasonable decision to treat past cumulative forecast errors, in particular relative cumulative errors (the reason is clarified later), as random variables that follow a similar distribution in the future seasons to the past ones. Similar decisions were taken by Fisher and Raman, 1996 in the famous Sport Obermeyer case study and by Cachon and Terwiesch, 2012 in their O'Neill case study.

Due to the lack of sufficient historical data, it is impossible to reliably estimate the distribution of forecast errors for individual products. We overcome this problem by working with relative cumulative forecast errors instead of absolute ones, and assuming that the relative cumulative forecast errors of different products for a given month are independently and identically distributed random variables (a similar approach is used by Cachon and Terwiesch, 2012). This enables us to pool historical cumulative forecast errors data over products to obtain enough data points to characterize a cumulative error, and therefore a demand distribution for each product-period combination. However, in the setting under consideration, we cannot test the validity of this assumption from the data itself, again because of the limited data available. Thus, we discussed the validity of this assumption with the company's experts and tested the sensitivity of the results of the model to a misspecification of the demand distribution. We discuss this robustness check in more detail in Section 5.4.

Mathematically, we derive cumulative demand distributions from relative cumulative forecast error distributions as follows: For each data point available, i.e. each season j of the available J years of history and each product $k = 1, \dots, K$, we define the relative cumulative forecast error, $RCFE$, of a cumulative forecast made at the start of the planning horizon for month t, which we denote by f_{tjk}, as

$$RCFE_{tjk} = \frac{f_{tjk} - a_{tjk}}{f_{tjk}}, \tag{107}$$

where a_{tjk} is the cumulative actual demand of product k in season j up to month t. It should be noted that the measure in equation (107) is undefined for $f_{tjk} = 0$. To solve this issue, we simply exclude the (very few) data points for which this happens from consideration. Combining all products and past seasons, we obtain a set of $P = K \times J$ $RCFE$ data points for each month $t \in \mathcal{T}$: the P points in this set, with element p denoted by $RCFE_{pt}$, form the empirical distribution of the random variable $RCFE_t$, from which the empirical probability distribution of the cumulative demand random variable CD_{kt} is readily obtained. Accordingly, the actual cumulative demand forecast at the start of the planning horizon for product k in month t, f_{kt}, the points $p = 1, \dots, P$ of the empirical distribution of CD_{kt}, denoted by CD_{pkt}, are obtained from:

$$CD_{pkt} = f_{kt} \cdot (1 - RCFE_{pt}). \tag{108}$$

This probability distribution is not only used to derive the scenarios of the stochastic model presented in the previous section, but also to evaluate the expected performance of all models used in this chapter's analysis.

Finally, given the probability distribution of demand obtained as previously explained, we use descriptive sampling (DS) to derive scenarios for the SGLSP-CDS model, as in Chapter 4, and we choose to use 10 scenarios, because this proved to be a valid compromise between the quality of the solutions and the computational complexity of the model. The S scenarios of the cumulative demand for each product k and period t combination are derived as:

$$CD_{kts} = F^{-1}_{CD_{kt}}\Big(\frac{s-0.5}{S}\Big) \qquad\qquad s = 1,\dots,S. \tag{109}$$

Because we sample from a discrete distribution, in case there is no point in the empirical distribution which exactly satisfies equation (109) for some s, we set the value of the cumulative demand in scenario s equal to the smallest point CD_{pkt} for which $F^{-1}_{CD_{kt}}(CD_{pkt}) \geq (\frac{s-0.5}{S})$.

It is worth mentioning that, for some products, the above scenario-generation process could result in negative expected demands for individual periods, because cumulative forecast errors data are pooled across all products. However, this does not affect the implementation of the CDS approach in our setting.

Additionally, we use the actual demand and forecast data of Seasons 2 and 3, along with that of other past seasons, to derive the demand distribution which is utilized both as input in the SGLSP-CDS model and to obtain the expected performance of all the models in the same seasons. In other words, for Seasons 2 and 3 we use an in-sample evaluation approach for all models, whereas for Season 1 we use an out-of-sample evaluation approach. We make this pragmatic choice because of the limited amount of data available.

5.4 Results

After deriving the problem's relevant input as explained in Section 5.3, we now conduct an analysis to reach the research objectives described at the start of the chapter.

Prior to showing the results, we make some general remarks which are relevant for the entire section. The different analyses presented involve solving one or more optimization problems, i.e. the SGLSP-CDS model and others which will be newly introduced shortly. All these problems are solved on a machine with a 2.30 GHz CPU and 64 GB of RAM using Gurobi 8.1 with a time limit of one hour, which, in most cases, leads to an optimality gap ranging from 1% to 3.5%. As discussed in the previous section, we use the actual data of three past seasons to evaluate the resulting solutions and draw our conclusions. For all seasons, the expected performance of the models' optimal plans is computed according to the same cumulative demand distribution which is used to derive the scenarios in the SGLSP-CDS model. In addition, we evaluate the models' performance for the actual demand realizations in the given season. For all analyses presented, we show and explain the results in detail for Season 1, and present additional results for Seasons 2 and 3 in Appendix B. Additionally, to maintain confidentiality, we normalize the sensitive original input and output quantities of the analyses.

5.4.1 Inventory drivers and their importance

In the first step of our analysis, we attempt to understand the impact of uncertainty as an inventory driver in the problem setting under consideration by comparing it to the other inventory drivers. We start by identifying and categorizing the most relevant inventory drivers as follows:

- Transportation time. This is an unavoidable time during which the produced products remain in inventory before they reach their sales region and can be sold to customers. As mentioned, products must be available in the sales region one month in advance, therefore, the production quantity in a given month stays in inventory for at least one month.
- Seasonal demand and limited capacity. The seasonality of demand is one of the most distinctive characteristics of the industry in which the firm operates. Due to the limited production capacity, this forces the firm to produce ahead of the peak demand season and use inventory in order to avoid unmet demand.
- Economies of scale. The agrochemical production processes are characterized by long, sequence-dependent setup times and minimum batch constraints, which make batch production necessary to avoid an excessive number of costly changeovers.
- Uncertain demand and desired service level. As discussed, the demand for CPPs can exhibit a high degree of uncertainty. Inventory is needed to buffer against this uncertainty to ensure that the desired service level is attained with the minimum financial investment in inventory.

To quantify the impact of the above drivers, we follow a stepwise approach. Focusing our attention on Season 1, in each step (which we call phase) of the approach we solve an optimization problem which determines the production plan that minimizes inventory costs considering a subset of inventory drivers. Specifically, following the order in which they are presented in the earlier provided list, we add an additional inventory driver in each phase and compare the resulting minimum inventory costs in the actual realized scenario of Season 1. The increase in the minimum inventory costs from one phase to the next serves as an indicator of the importance of the newly added inventory driver. To avoid unnecessary repetition, we present only the mathematical formulation of the (deterministic) problem of Phase 3 in Appendix C. The problems for Phases 1 and 2 can be readily derived from that by deleting or modifying some constraints, as explained in subsequent sections.

In Phase 1, we consider only the transportation time as an inventory driver. We assume that demand is known and equal to the realized demand in the season. We solve a simple, linear deterministic production planning problem with the objective of minimizing total inventory costs subject to meeting the demand and considering the lead time. As in the SGLSP-CDS model, we treat initial inventory as a variable to make the analysis independent of the starting condition. We also force the ending inventory to be at least as large as the initial one, thus creating a cyclical plan. Capacity constraints and the production process-related constraints of the standard deterministic GLSP are disregarded because they are associated with other inventory drivers not considered at this stage. Formally, we solve the model shown in Appendix C excluding constraints (112)-(116) and (121)-(122).

In Phase 2, we consider additionally the limited capacity. The linear deterministic model used for Phase 1, which uses seasonal demand as input, is extended by the addition of simple capacity constraints that ensure that the total production time in each month does not exceed the available capacity. These constraints are obtained by excluding the second term on the left-hand side of constraints (112) in Appendix C.

In Phase 3, we also consider economies of scale. As in the SGLSP-CDS model, we do that by considering sequence-dependent setup times and minimum batch constraints. The deterministic model of the previous phase is extended by adding the standard production-process-related constraints of the deterministic GLSP. The complete model formulation is presented in Appendix C.

In Phase 4, we consider demand uncertainty and the target service level. To that end, as opposed to the previous three phases, we solve the stochastic SGLSP-CDS model presented in Section 5.2. The optimal plan of the SGLSP-CDS model, and thus the resulting actual inventory costs in Season 1, is a function of the δ service-level target. In order to estimate the effect of varying this parameter, we solve the stochastic model multiple times with different target δ service levels. We show the results for δ service-level targets equal to 0.96, 0.98 and 0.99. The choice of these relatively high service-level targets is motivated by the fact that customers expect high service levels and that there are substitutes of the products in the market, which drive AC to aim for such targets.

The optimal monthly inventory costs of the different phases are plotted in Figure 30.

Focusing our attention on the inventory cost curves of the first three phases, we see that the differences between them are really small. Transportation times cause significant unavoidable inventory costs, whereas the addition of seasonal demand and limited capacity, as well as economies of scale, do not cause a significant increase in inventory costs. A comparison of the curve of Phase 3 to the curves of phase 4 shows a significant increase in inventory costs, which clearly indicates that the uncertainty of demand and the desired service-level targets is an inventory driver which significantly amplifies the impact of the other drivers. This effect of demand uncertainty emphasizes the importance of appropriately considering this inventory driver when planning. For the company, this means that it is valuable to invest time and resources to improve their planning under uncertainty

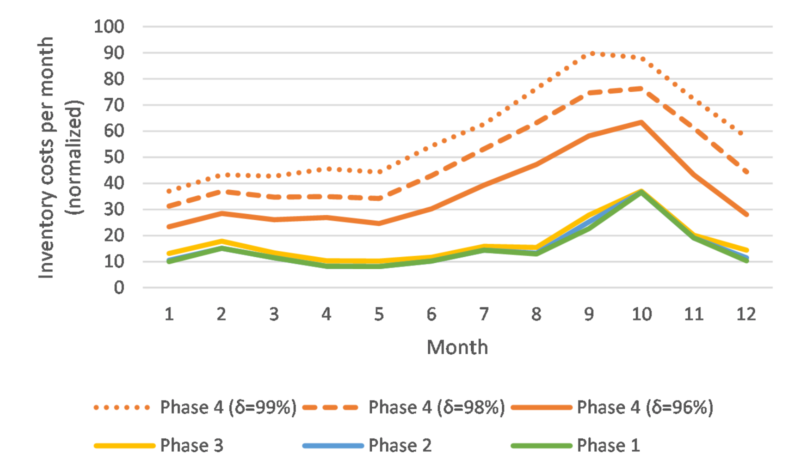

Figure 30: Comparison of inventory drivers – season 1

and to improve the accuracy of their demand forecasts. In addition, the results show that a relatively small increase in the target service level requires a significant increase in the necessary inventory investment, emphasizing the importance of carefully choosing an appropriate target.

5.4.2 Inventory-service level trade-off and assessment of the current planning procedure

Based on these results, in the next step of the analysis, we focus on understanding the trade-offs between inventory costs and service level in detail, as well as on comparing the performance of the developed stochastic model and the planning process currently used by the plant. Conceptually, the latter works as follows. First, forecasts for the next season (and beyond) are generated in an S&OP meeting. Second, these forecasts are adjusted, in nearly all cases inflated, by individual countries' sales experts considering the company's goals and the uncertainty of demand in their region. In the final step, a production plan to meet the inflated forecasts is developed. Essentially, in this sequential process, safety stocks are determined by experts, in the form of inflated forecasts. Subsequently, a deterministic planning problem is solved to meet the resulting safety stock targets. This type of approach, in which demand uncertainty is dealt with before the planning problem is solved, is frequently used in practice and in the literature (Boulaskil et al., 2009), despite it being considered suboptimal. In the specific case of AC, two main issues of the approach can be identified. The first is its sequential nature, which makes it impossible to consider the interaction and possible synergies between safety and other types of stock. The second is the unsystematic representation of uncertainty and the lack of clear rules for determining the required safety stocks. Indeed, no systematic statistical analysis is conducted to characterize demand uncertainty, and experts decide safety stock levels relatively subjectively.

In our analysis, we replicate the company's current planning approach as follows: to obtain the company's planning decision for a particular season, we solve a deterministic GLSP which aims at finding the production plan which minimizes inventory costs while satisfying a demand equal to the inflated forecasts that are obtained in the second step of the company's planning procedure (we refer to this model as the *GLSP-IF model*). The mathematical formulation of the model is identical to that of the Phase3 Model shown in Appendix C, with the exception that the parameter CD_{kt} now represents the cumulative demand for product k up to period t obtained from the experts' inflated forecasts at the start of the planning horizon. In the remainder of this chapter, we will use the terms GLSP-IF and SGLSP-CDS to refer to the inventory planning approaches, and GLSP-IF *model* and SGLSP-CDS *model* to refer to the mathematical model used to solve the inventory problem in the corresponding approach. A similar naming strategy is applied to all other approaches introduced in this section.

For each analyzed season, we first obtain the production plans from the SGLSP-CDS model and from the deterministic GLSP-IF model. Because the optimal plan of the SGLSP-CDS model is a function of the service-level target, we solve the stochastic model multiple times for different target δ service levels. This allows us to show the trade-off between service level and inventory costs inherent to the problem, and to obtain a more insightful comparison of the SGLSP-CDS and GLSP-IF. The optimal plan of the GLSP-IF is instead unique, because the service level is entirely determined by the inflated forecasts which serve as input to the model. Thereafter, we simulate and compare the performance of the approaches, both in expectation, i.e. assuming that the true demand distribution is the one used to derive the scenarios for the SGLSP-CDS models, and in the actual scenario.

For Season 1, the results in expectation are shown in Figure 31. As can be seen, the plan developed by the SGLSP-CDS outperform the one obtained by using the GLSP-IF in expectation. Specifically, it can be concluded that approximately 30% of inventory costs can be saved while keeping the δ service level constant by using the SGLSP-CDS instead of the GLSP-IF. Alternatively, keeping inventory costs constant, the service level can be improved by approximately 4% using the stochastic model for planning.

The results of the comparison for the actual demand in Season 1 are shown in Figure 32. As can be seen, in the actual scenario, the SGLSP-CDS only slightly outperforms the GLSP-IF. This suggests that the actual scenario's demand realization is a favorable one for the GLSP-IF.

Similar analyses for Seasons 2 and 3 are shown in Figures 42-45 in Appendix B. In these two seasons, the SGLSP-CDS clearly outperforms the GLSP-IF, both in expectation and in the actual scenario.

The above analysis relates to both of our research objectives. First, it clearly shows the

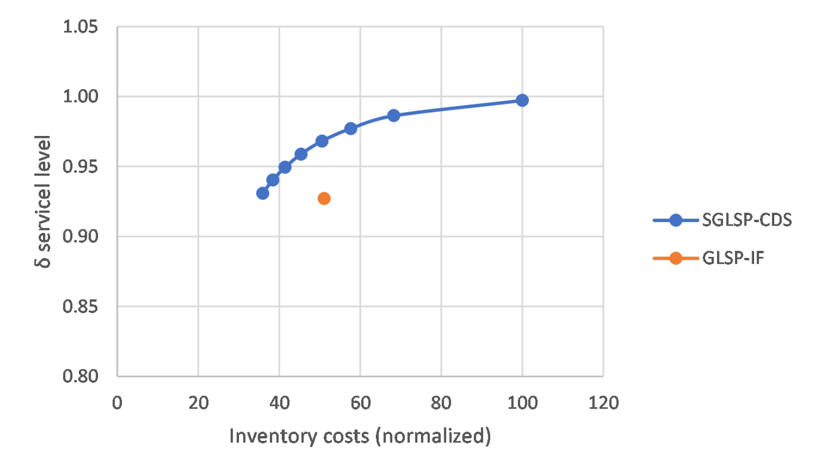

Figure 31: Expected performance of the SGLSP-CDS and GLSP-IF – Season 1

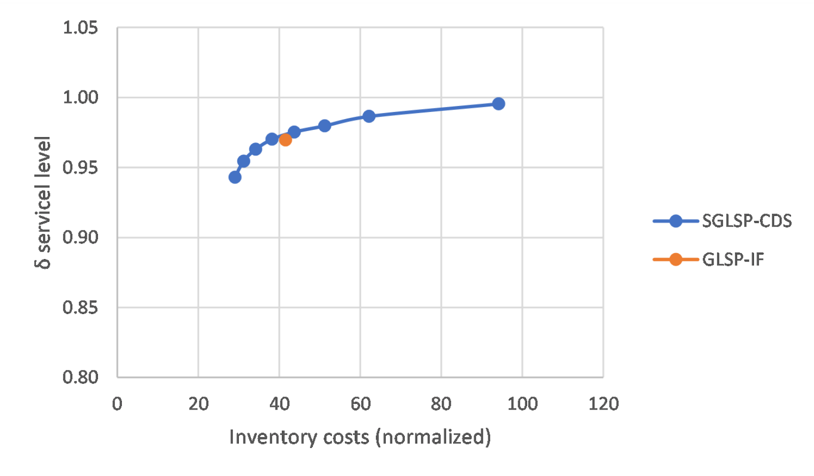

Figure 32: Performance of the SGLSP-CDS and GLSP-IF models – Season 1

level of investment in inventory that is needed to achieve a certain service level. The classical trade-off between the cost minimization and service-level maximization objectives provides practitioners with a useful tool for making, justifying and communicating inventory management decisions. Second, it shows that the newly developed SGLSP-CDS can deal with uncertainty more effectively than the planning approach currently used.

5.4.3 Comparison of inventory plans of the newly developed and the existing planning approaches

We now direct our attention to understanding the origins of the observed performance difference between the two approaches to inventory planning by first looking at how their optimal planning decisions differ. To this end, we compare the optimal plan of the GLSP-IF model with the optimal plan of the specific SGLSP-CDS model that achieves the same expected aggregate service level (according to the demand distribution implied by the scenarios used in the stochastic model) but with lower expected inventory costs in Season

1. In other words, we investigate the differences in the optimal planning decisions of the GLSP-IF and the SGLSP-CDS models, assuming that the aggregate target service level of the company is the one reached in expectation by the GLSP-IF plan. First, although both plans achieve the same aggregate service level, the service level achieved in expectation for the individual products differ, and this could explain the difference in expected holding costs. To quantify the extent to which this aspect can explain the difference in the inventory costs performance of the two approaches, we solve the SGLSP-CDS problem with individual product target service levels set equal to those achieved in expectation by the GLSP-IF plan. Comparing the resulting expected inventory costs to those of the original SGLSP-CDS plan, we find that of the 30% savings in expected inventory costs, 9% stem from the difference in the individual products' service levels achieved in expectation. Therefore, a large part of the difference in expected holding costs between the SGLSP-CDS and GLSP-IF plans remains unexplained. The remaining important difference between the two plans is therefore the production schedule. Comparing the two production schedules, we find that, although the number of changeovers and the overall utilization of the production resources is almost identical, the decisions on the timing and size of the production lots differ substantially. We show the two aggregate plans in Figure 33.

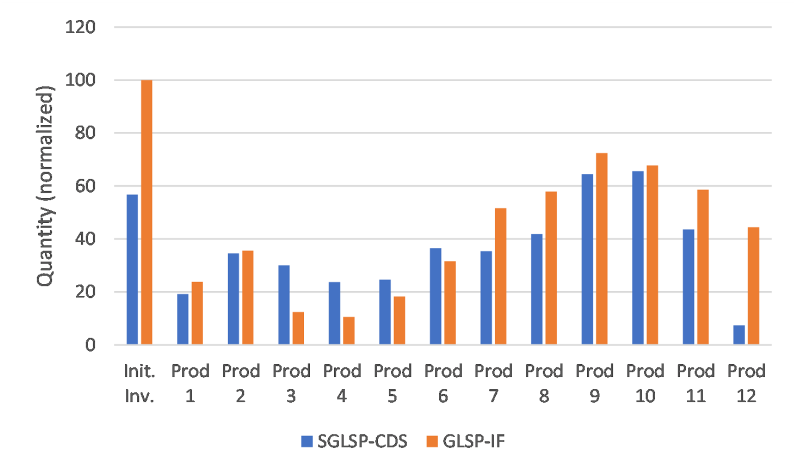

Figure 33: Comparison of resulting production plans between SGLSP-CDS and GLSP-IF models – Season 1

According to the figure, the most noticeable difference is that the GLSP-IF plan starts with a much larger inventory than the SGLSP-CDS plan. As previously mentioned, the difference between the two plans is ultimately due to the way the two approaches treat demand uncertainty and calculate the optimal safety stocks. In the SGLSP-CDS, safety stocks are calculated endogenously by the planning model, whereas in the SGLSP-IF they are calculated exogenously and then input into the planning model in the form of inflated demand forecasts. Therefore, the underlying input demand to the GLSP-IF model explains the difference between the production plans and thus the inventory costs of the approaches. In Figure 34, we show the total demand for each period of the planning horizon that the GLSP-IF model aims to fulfill, which is a result of a demand estimation process and safety stock decisions, and the expected total demand for each period which is assumed by the SGLSP-CDS model when simultaneously determining the production schedule and the safety stock.

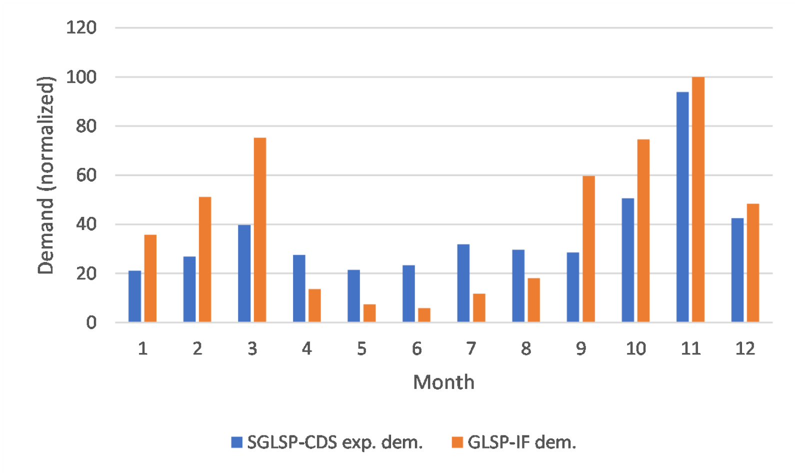

Figure 34: Comparison of input demand between SGLSP-CDS and GLSP-IF models – Season 1

The graph clearly shows that the GLSP-IF plan starts with higher inventory because it aims to build substantial safety stocks in the first periods of the off-season, when demand is rather low. For Seasons 2 and 3, however, this difference in demand in the first months is found to be less striking (the analysis is shown in Figures 46-49 in Appendix B); as a result, also the optimal production plans do not differ as much as for Season 1. In those seasons, however, the expected savings of the SGLSP-CDS over the GLSP-IF are also lower, although still significant.

This analysis confirms that the previously observed advantage of the SGLSP-CDS over the GLSP-IF originates from the way the former approach determines safety stocks to buffer against demand uncertainty. Conceptually, this advantage comes from two sources. One is the use of available statistical information on demand uncertainty in the form of demand distributions instead of a subjective estimation of the demand uncertainty, which is used when generating inflated forecasts. The other is the use of a stochastic planning model, which plans safety stock and production in an integrated fashion, i.e. simultaneously, as opposed to the currently used planning approach, where the safety stocks are determined before the planning model is solved. When determining the safety stock, the sequential approach of the GLSP-IF fails to consider many features of the problem, such as limited capacity and setup times. In other words, contrary to the SGLSP-CDS, the current sequential approach disregards the fact that the optimal safety stock is a function of (depends on) the production plan.

Given these two sources of the advantage of the SGLSP-CDS over the SGLSP-IF, it is insightful to determine which source leads to the most significant improvement over the current planning procedure. Accordingly, we investigate the potential benefits of using a statistically-based but sequential planning process for determining the safety stocks over the currently used method of forecast inflation.

In the first step of the current approach used by the company, forecasts for each article and period are inflated subjectively, leading to safety stock targets. A different approach to determining safety stocks is to perform this step using the statistical knowledge of the demand process which we use for deriving scenarios in the SGLSP-CDS model. Specifically, we consider the following planning approach:

1. Derivation of the relevant probability distribution of demand. In this step, the cumulative demand distributions for each product-period pair are estimated, as for the SGLSP-CDS model.

2. Safety stock calculation as a demand percentile. Here, a certain percentile of demand for each product-period combination is determined based on the service-level target. We assume that the same service-level target, and thus the same percentile, is used for all products and periods, although it could potentially be different for each combination.

3. Production planning. We solve a deterministic GLSP aiming at minimizing inventory costs subject to meeting the production targets corresponding to the percentiles of the demand distributions determined in step 2. Specifically, the mathematical model is identical to that shown in Appendix C, with the exception that the parameter CD_{kt} now represents the percentile of the cumulative demand distribution for product k up to period t determined in step 2.

The second step is the statistically-based equivalent to the currently used forecast-inflation method. We solve the inventory problem using this planning approach, which we denote by GLSP-SIF (where SIF stands for statistically inflated forecasts), with different percentile targets in the second step, and compare the results with the ones of the GLSP-IF and SGLSP-CDS. The results of the analysis for Season 1 are shown in Figures 35 and 36.

Each point in the GLSP-SIF curve corresponds to a solution obtained using a different percentile of the demand distribution for all product-period combinations to determine the production targets. Specifically, from left to right the targets percentiles are 70%, 72.5%, 75%, 77.5% and 80%. As can be seen, in expectation, the GLSP-SIF performs better than the GLSP-IF, showing that it is beneficial to use the available statistical knowledge of the demand process, even when using a simple rule of thumb to make safety stock decisions in a sequential planning approach. Also, predictably, the GLSP-SIF is outperformed by the GLSP-CDS, because the latter optimizes the safety stock and planning decisions simultaneously, instead of sequentially. The comparison between

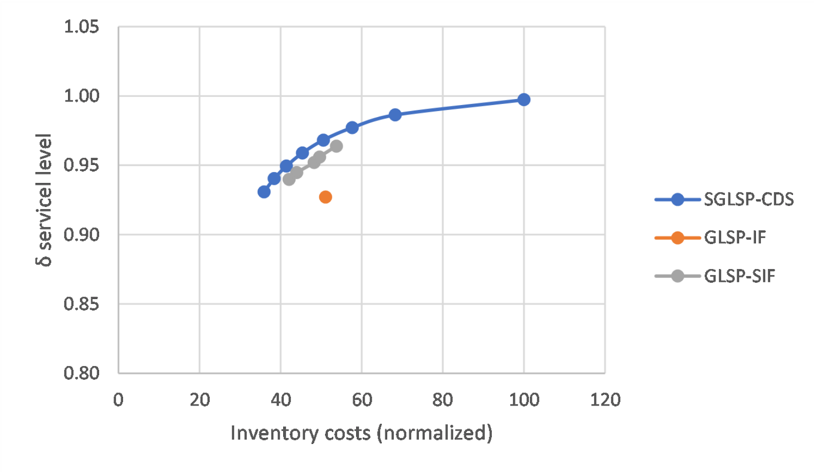

Figure 35: Comparison of SGLSP-CDS, GLSP-IF and GLSP-SIF models – expected performance in Season 1

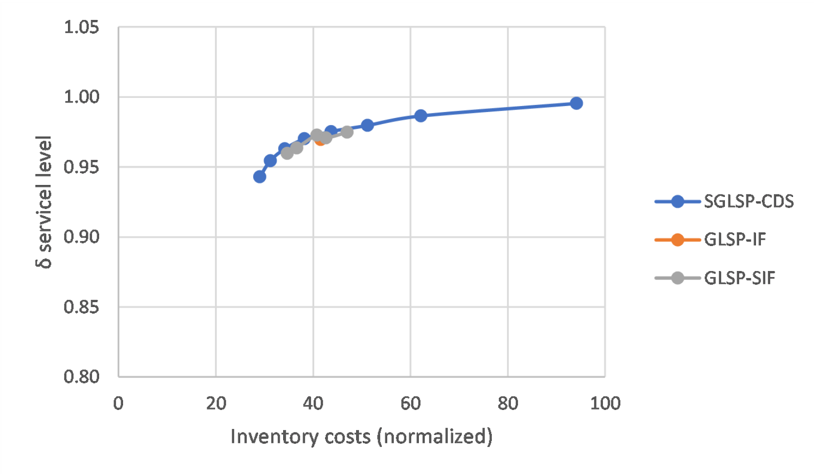

Figure 36: Comparison of SGLSP-CDS, GLSP-IF and GLSP-SIF models – actual performance in Season 1

the expected performance of the three models provides additional information on the shortcomings of the currently used planning approach. As can be seen, the difference in performance between the SGLSP-CDS and the GLSP-SIF is smaller than that between the GLSP-SIF and the GLSP-IF. This shows that the major drawback of the current approach is the lack of a systematic, statistically-based method to determine safety stock levels. Consequently, implementing the proposed GLSP-SIF, even if the sequential nature of the planning approach is kept, would benefit AC. However, it is also important to note that in the GLSP-SIF there is not a clear relationship between the percentile of demand used to determine safety stocks and the expected achieved service level, thus complicating the application of the approach in practice. For example, choosing the 75th percentile does not lead to a 75% service level. Also, independent of the expected service level achieved, the resulting plan is not necessarily the cost-optimal way to reach the target, because the interaction between periods and products are disregarded in this sequential approach, since the safety stock targets are determined independently for each product-period combination. Additionally, due to the EOH constraints, large percentiles may lead to infeasibility of the problem. Finally, we see that it would be beneficial for the company to take into consideration switching from a sequential planning approach to an integrated one, in which safety stocks and production decisions are optimized simultaneously.

The difference between the GLSP-SIF and the other approaches in the actual scenario of Season 1 is not clear-cut, but it seems that, as it is the case in expectation, the GLSP-SIF slightly outperforms the GLSP-IF and is marginally outperformed by the SGSLP-CDS. Similar results are obtained for Seasons 2 and 3, with the corresponding graphs shown in Figures 50-53 in Appendix B. However, in these seasons, the relationship between the different models in the actual scenario is much clearer, and comparable to the relationship between the models in expectation in all seasons.

5.4.4 Robustness of the SGLSP-CDS to misspecifications of the demand distribution

In Section 5.3 we explained that we obtained the results shown so far assuming that the relative cumulative forecast errors of all products are iid. We now try to test the robustness of the SGLSP-CDS model to a different distribution assumption.

After a discussion with the company's experts, we obtain the different distribution assumption for the robustness check as follows: We use a classification scheme developed by the company for forecasting purposes to group the products into three clusters. We then assume that the relative cumulative forecast errors of all the products in each cluster are iid and thus derive a different demand distribution for all products by repeating the procedure presented in Section 5.3 for each cluster.

Under the resulting new demand distribution assumption, we solve the SGLSP-CDS model again with different service-level targets to obtain the new optimal inventory and production plans. Subsequently, we compare the expected performance of these plans, according to the newly derived demand probability distribution, with that of the original plans obtained by the SGLSP-CDS model without clustering and the plan of the GLSP-IF model. We therefore evaluate the expected performance of all models under the assumption that the actual demand distribution is the one obtained using the clustering of products. The results of this analysis are shown in Figure 37.

As can be seen, the SGLSP-CDS model with clustering (denoted by "clustering" in the graph) outperforms not only the GLSP-IF, but also the SGLSP-CDS model without clustering (denoted by "no clustering" in the graph). This is to be expected, because to the fact that the plan of the SGLSP-CDS model with clustering is optimized for the the demand distribution implied by the clustering, which is also used to obtain the expected performance of the all plans shown in Figure 37. However, it is important to note that

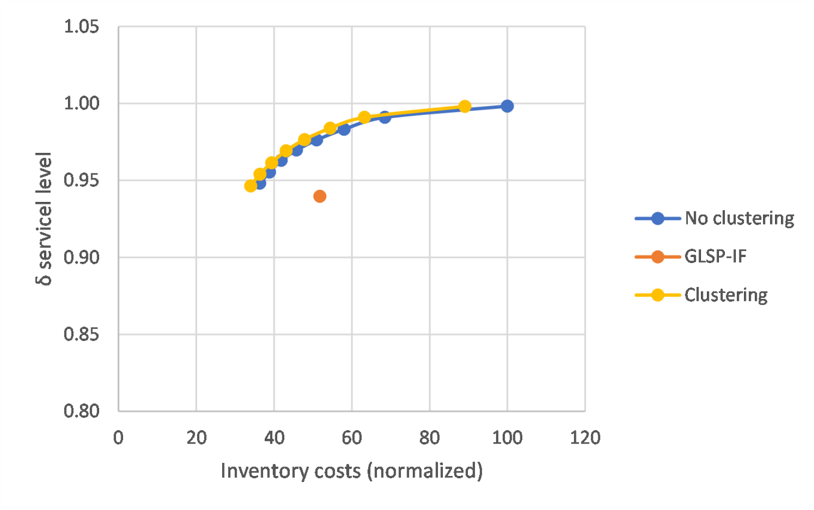

Figure 37: Robustness check to misspecifications of the probability distribution of demand – Season 1

the performance of the SGLSP-CDS model without clustering is very close to that of the model with clustering, suggesting that the plans of the stochastic model are indeed robust to misspecifications of the probability distribution of demand. We also observe that the SGLSP-CDS model without clustering still clearly outperforms the GLSP-IF. This appears to indicate that even if the developed stochastic planning model uses an incorrect demand probability distribution assumption, it still performs better than the current approach. This strengthens the conclusion that the SGLSP-CDS model provides benefits over the currently used planning approach. The robustness of the stochastic model is further confirmed by the results of the same analysis in Seasons 2 and 3, which are shown in Figures 54-55 in Appendix B. In the remainder of this chapter, we will thus continue using the demand distribution obtained without clustering the products to show the results of our analysis.

5.4.5 Aggregate versus individual service-level targets

In the analysis presented so far in this chapter, we have used aggregate service-level constraints, meaning that the service level is defined and controlled as aggregated over all products. We made this choice because it simplified the process of obtaining and discussing the results of our analysis with our research partner. Theoretically, however, service-level constraints can be disaggregated (aggregated) along many dimensions, such as product, time, customer, or geographical region. Because the choice of the aggregation level of the service-level constraints has important consequences on the required inventory investment to reach a certain target, we now analyze the effects of disaggregating the service-level constraints for individual products. Specifically, we enforce a δ service-level constraint for each individual product. Although we focus on this specific relevant dimension, we note that the analysis can be repeated for others, providing further insights to the company to help it setting the right service-level targets, which is a highly challenging task in practice. In principle, the targets can also differ between products, however, in our analysis we assume that they are identical for the sake of simplicity.

We show the expected performance of the SGLSP-CDS with individual service-level targets (which we denote by SGLSP-CDS-ind) in Season 1 in Figure 38. As for the previous analyses, the vertical axis represents the expected aggregate service-level performance of the different plans. By comparison, in the same graph we again show the performance of the SGLSP-CDS with aggregate service-level targets and of the GLSP-IF.

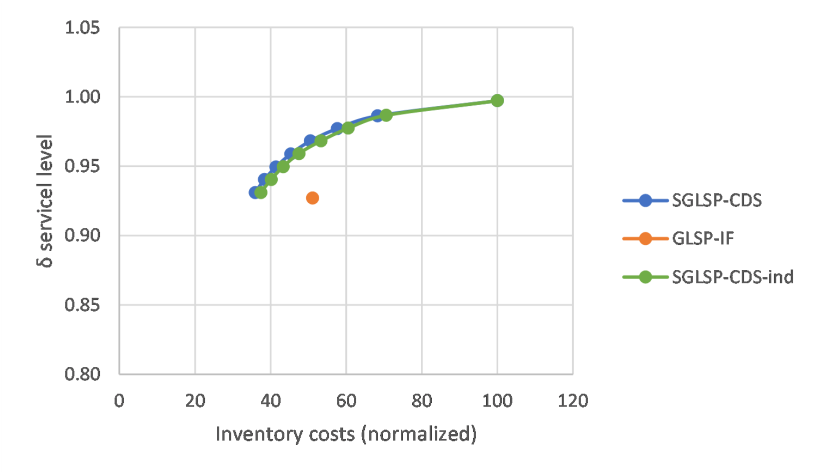

Figure 38: Expected performance of the SGLSP-CDS-ind, SGLSP-CDS and GLSP-IF models – Season 1

Predictably, it is shown that the expected costs necessary to achieve any given target δ service level are higher if this target must be met for each individual product instead of over all products. A reason for this is that using an aggregate service-level constraint allows the company to overachieve the target for the less expensive (in terms of inventory costs) products and underachieve it for the more expensive ones, while still meeting the overall target. Another reason is that, even if all products were equally expensive, a bad service-level performance for one product in one scenario can be compensated for by a good service-level performance for another product in the same scenario, thus allowing risk-pooling. In Figure 39, we show the relationship between the optimal service level chosen by the SGLSP-CDS and the unit inventory costs of all products when the target δ service level is set to 98%. As mentioned, the optimal decision is to overachieve the service-level target for products with low inventory costs and underachieve it for those with higher inventory costs.

Although Figure 38 shows that there is a difference in the inventory cost performance between the SGLSP-CDS and the SGLSP-CDS-ind, this difference is small. This can largely be attributable to the consideration of high service-level targets in our analysis. These high targets leave little room to the SGLSP-CDS for overachieving the target for the cheaper products, which therefore cannot compensate for a large underachievement

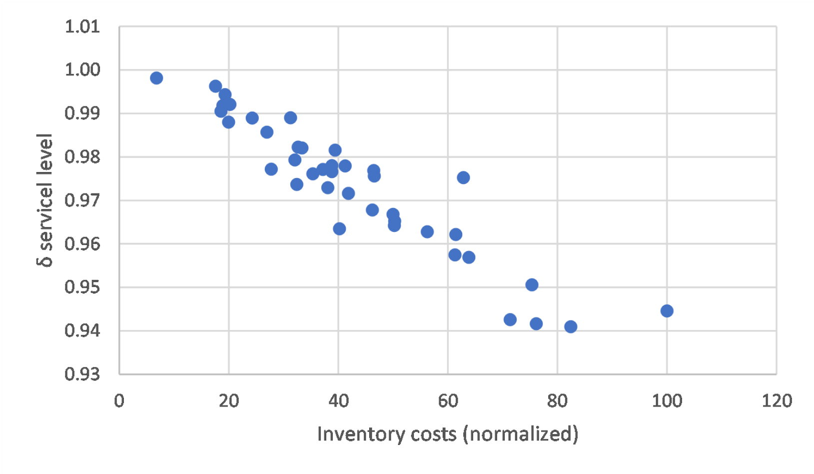

Figure 39: Optimal service level – inventory costs relationship with aggregate service level $\delta = 98\%$ – Season 1

of the target for more expensive products. Indeed, in Figure 39 we see that in order to reach an aggregate target of 98% it is not possible to aim for a service level lower than 94% even for the most expensive products without increasing total inventory costs. To summarize, this analysis shows that in the problem setting under consideration, given that high service levels are targeted, AC would not need to substantially increase its inventory investment if it decided to aim for a comparable service level for each individual product. It also shows that using an aggregate service-level constraint does not result in an excessively uneven distribution of service levels among products.

5.4.6 Cost-based model to support service-level target setting

We conclude our study by discussing the problem of setting the right target for the service-level constraints. Thus far, we have shown the effects on inventory costs of changing the target service level by solving the SGLSP-CDS (or SGLSP-CDS-ind) model for different δ service-level targets, but have not discussed how AC should choose the right target (or targets, depending on the aggregation level of the service-level constraints).

Due to the uncertainty of demand, each possible plan implemented by AC implies two risks: on the one hand the risk of experiencing a shortage, and on the other hand that of carrying unnecessary costly inventory. In theory and in practice, the most common ways to manage these risks are two. One is to enforce service-level constraints, the method applied so far in this chapter, which requires decision makers to set the target service levels. The other is to use shortage costs: once these costs are defined, they are added to the objective function and traded-off against inventory costs. Using the latter method, the service-level targets of different products are essentially decided by the products' underage and overage costs, instead of by the decision maker. The target service levels chosen by the latter *cost-based* approach can be used in practice as a starting point for setting the right targets for the products in the former *service-level-based* approach.

Therefore, we now develop a cost-based model of the mid-term production and inventory planning problem faced by AC, with the aim to help them setting the right service-level target and thus determining the necessary inventory investment to reach their goals.

In the lot-sizing literature, it is common to control shortages by penalizing each backordered unit by a certain cost in the objective function. However, there are other ways of "costing" a shortage. Other common ways include, for example, defining a fixed cost per stockout occasion or a fractional charge per unit short per unit time (see Silver et al., 2016 for a more in-depth discussion on different methods of costing shortages). In the developed cost-based model, we control stockouts by introducing a penalty for each unit of backlog per time period. We derive the cost-based model from the SGLSP-CDS model by modifying the objective function and eliminating the service-level constraints of the latter. Specifically, letting p_k be the penalty cost per period of one unit of backlog of product k, we formulate the cost-based model, named SGLSP-CDS-p, as follows:

SGLSP-CDS-p Model

$$\text{Min} \sum_{k\in\mathcal{K}}\sum_{t\in\mathcal{T}} h_k \cdot \Big(\sum_{s\in\mathcal{S}} pr_s \cdot IP_{kts}\Big) + \sum_{k\in\mathcal{K}}\sum_{n\in\mathcal{N}}\sum_{m\in\mathcal{M}} h_k \cdot x_{knm} + \sum_{k\in\mathcal{K}}\sum_{t\in\mathcal{T}} p_k \cdot \Big(\sum_{s\in\mathcal{S}} pr_s \cdot BL_{kts}\Big) \tag{110}$$

Subject to constraints (92)-(100), and (102)-(106).

Because it is difficult, if not impossible, to correctly estimate the penalty costs p_k in practice, we solve the problem for different values of this parameter. To this end, we assume that the backlog costs of all products are directly proportional to their profit margins. Furthermore, we obtain different values of p_k by varying the proportionality factor, which we denote by q. We present the results of the analysis in Figure 40, showing the expected performance of the SGLSP-CDS-p, as well as that of the SGLSP-CDS and GLSP-IF by comparison, in Season 1.

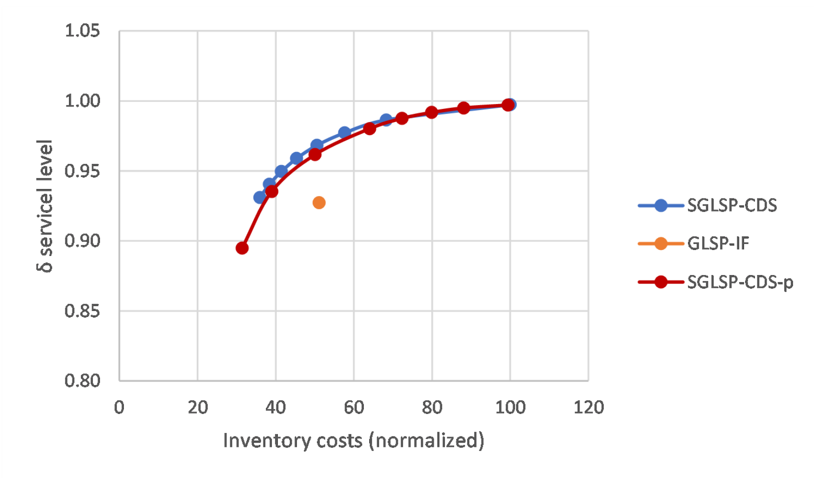

Figure 40: SGLSP-CDS-p – Expected performance, Season 1

From left to right, the points in the SGLSP-CDS-p curve correspond to an increasing value of the constant of proportionality q used to obtain the penalty terms (the specific values are confidential, and hence, not shown). As expected, as the backlog costs increase, the expected inventory costs and aggregate service level also increase, because the SGLSP-CDS-p model accepts higher inventory costs in order to avoid the relatively more expensive backlogs. Moreover, importantly, these results give decision makers at AC initial valuable insights into the problem of setting the right service-level targets. Indeed, given the company's estimate of the constant of proportionality q, and thus backlog costs, the graph shows the service-level target δ for which they should aim. Of course, however, solving the SGLSP-CDS model with an aggregate target δ we will not obtain the same optimal production/inventory plan as when solving the corresponding SGLSP-CDS-p model with the constant of proportionality q that leads to the solution with an aggregate service level of δ. Indeed, as can be seen by comparing the SGLSP-CDS and the SGLSP-CDS-p curves, the cost-based model prescribes higher inventory costs than the service-level-based model. Specifically, the optimal plans of the two models differ in the service-level targets set for the individual products. To see how, for illustrative purposes, we compare the optimal service level of all products chosen by the SGLSP-CDS model with a target of $\delta = 98\%$ to that chosen by the SGLSP-CDS-p model which leads to the same aggregate service level. In Figure 41, we plot the difference between the optimal δ service level chosen by these two models for all products against their profit margins.

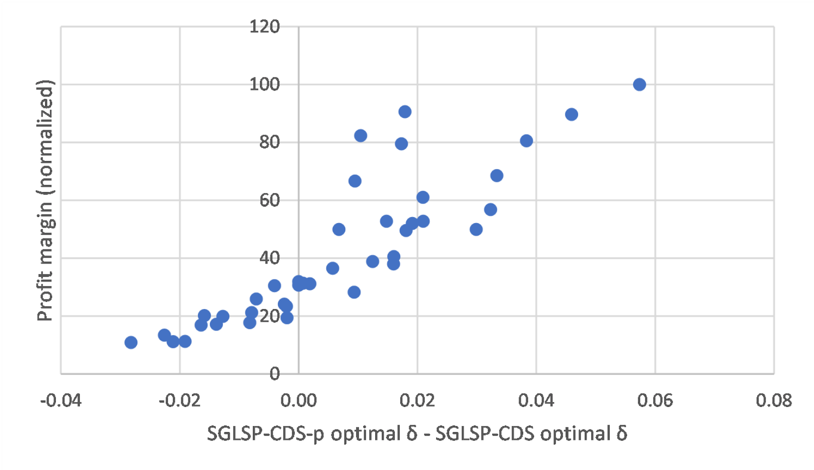

Figure 41: Difference in the optimal δ service level between SGLSP-CDS and SGLSP-CDS-p models as a function of the products' profit margin – Season 1

As expected, compared to the SGLSP-CDS model, the SGLSP-CDS-p model chooses a higher service level for products with a large profit margin, although it achieves the same aggregate δ service level. This choice leads to larger inventory costs, but consequently also to lower backlog costs, and the decrease in the latter outweighs the increase in the former, thus minimizing the sum of the two costs. The company could, therefore, look at the service-level targets chosen by the SGLSP-CDS-p for a given q as a starting point to set individual targets for its products.

5.5 Conclusions

To summarize, in this chapter, we developed an integrated lot-sizing and scheduling model under uncertainty and used it to conduct an analysis to help AC quantify the impact of uncertainty in their mid-term inventory planning problem and make more effective and well-informed inventory decisions.

The proposed model, a stochastic version of the GLSP solved with the CDS approximation method presented and studied in Chapter 4, is able to cope efficiently with the major challenges faced by the company in the complex setting in which it operates, namely demand seasonality, demand timing and quantity uncertainty, and sequence-dependent setup times. Using this model, we conducted an extensive analysis which enabled us to provide insights into the company's planning problem. We also believe that the model can serve as a starting point to support actual future mid-term inventory planning.

We have shown that demand uncertainty is the most important inventory driver in the setting analyzed. Because of the seasonal nature of demand, uncertainty actually takes two forms: demand scale and demand timing uncertainty. As a result, it is crucial to manage both types of uncertainty appropriately when planning, and our proposed approach appears to be a promising way to do that. Moreover, this result also signals that it is beneficial for AC to invest in improving forecast accuracy.

The developed planning model allows to clearly visualize and quantify the trade-off between inventory costs and service level and can be used for evaluation and planning purposes. Indeed, performance evaluation and decision-making should become more transparent once the relationship between the two aforementioned key performance indicators is made clear and intuitive for all the stakeholders involved in the planning process.

Moreover, we simulated the performance of the stochastic model for three past seasons, and showed that it outperforms the currently used planning approach, both in expectation and in the actual scenario. In a further step, we also highlighted the shortcomings of the current planning method compared to the newly proposed planning approach.

Finally, we developed a cost-based stochastic lot-sizing and scheduling model and compared its behavior and performance to those of the initial service-level-based model. This new model represents an alternative approach to inventory management and can complement the originally proposed model by providing guidance in the process of selecting the right service-level targets for the company's products.

The nature of the market under consideration and the limited available data make it challenging to obtain an estimate of the probability distribution of demand. We have successfully conducted a limited test of the robustness the model to a misspecification of the demand distribution. However, future research could extend this analysis and test ways to improve the estimate of the demand distribution.

The performance of the model in past seasons could also be tested by applying the model under a rolling horizon planning approach, because that would provide a more accurate estimate of its performance. This would also facilitate applying the model in future planning efforts. Moreover, again for both evaluation and future planning purposes, the model could be extended to consider re-planning opportunities within the planning horizon, for example by using a static-dynamic or a dynamic uncertainty strategy as opposed to a static one. This different uncertainty strategy would provide superior solutions in terms of inventory costs, but might lead to planning nervousness, a significant increase in computational complexity of the problem and the need to use a solution approach other than the CDS. Indeed, with a static-dynamic or dynamic strategy, a multi-stage stochastic program is needed to model the problem, in which decisions in each stage are dependent not only on the past realization of demand and past decisions, but also on the conditional probability distribution of demand in the remaining periods of the planning horizon, thus also requiring an estimate of these distributions. As an alternative solution to transforming the model into a multi-stage stochastic program, the model can be extended by considering future re-planning opportunities under rolling horizon planning using the methodology introduced by Tavaghof-Gigloo and Minner (2020).

Another possible extension of the proposed model is to consider the bottling and labeling steps of the production process of CPPs at the plant to consider the capacity of these production stages more accurately. This could also be coupled with modeling the demand processes at the article level instead of the CPP level. Similarly, the model could be extended by considering the production decisions for AIs or their availability, which we assumed to be unlimited in our approach.

Finally, different scenario-generation methods could be evaluated, especially those able to provide scenarios which are paths through a scenario tree that respect demand dependencies. Coupled with the standard scenario-approximation methodology of stochastic lot-sizing problems presented in Chapter 4.4.1 (instead of the CDS one), this would allow to measure the variance of the objective function and thus control risk. However, these scenario-generation methods would require that the dependence between all the product-period combinations are first estimated and then generate scenarios which obey these dependencies. This process is especially challenging in the real-world setting analyzed, due to the number of products considered and the limited availability of demand data.

6 Conclusions

In this thesis, we studied the inventory management problem of a firm selling products with a seasonal and uncertain demand. Our focus was on simultaneously considering the quantity and timing uncertainty of demand and the related inventory decisions, a topic that has not been addressed in sufficient depth in the literature. When the products' selling seasons are stochastic, the right inventory timing decision is as important as the right inventory quantity to the success of the firm. Therefore, the correct representation of the timing uncertainty and the role of inventory timing in buffering against this uncertainty should be explicitly considered. In particular, we focused on showing the relationship between the problem's parameters and the inventory decisions of the firm, highlighting the (practical) relevance of the problem and developing tools to manage it in complex real-world settings.

We began this thesis by illustrating the analyzed inventory problem and its practical relevance with an in-depth description of the agrochemical industry and the inventory challenges faced by a firm manufacturing crop protection products (CPPs). Despite the degree and importance of the timing uncertainty of demand for these products, we showed that no published studies conducted in this industry directly address or propose solutions to the specific challenges posed by this key characteristic of demand.

In Chapter 3, we started our analysis of the inventory problem by investigating how the characteristics of the problem, i.e. the stochastic properties of the selling season of the product and the cost parameters of the product such as its profit margin and holding costs, influence the optimal inventory timing and quantity, and the profits of the firm. With this analysis, we filled a gap in the literature on inventory management of products with seasonal and uncertain demand. This literature has indeed paid almost no attention to demand timing uncertainty and the inventory timing decision, with the exception of Schlapp and Fleischmann (2020).

We first presented the model and summarized the analytical results of Schlapp and Fleischmann (2020), to comprehensively define the problem setting and characterize the optimal inventory timing and quantity of the firm, as well as to point out the main differences with the classic newsvendor model that this model is built on.

Then, we used this knowledge to conduct a numerical study aimed to answer our research questions. We first discussed all parameters varied across the problem instances solved in detail and justified our reasons for choosing the values of each parameter considered in our study. Given that there exists no known closed-form solution to the problem, we proposed a scenario approximation of the problem and a grid-search algorithm to solve it, to obtain good quality solutions in a reasonable time. We also defined a newsvendor-like naïve inventory policy commonly used in practice, which makes the product available as soon as possible to minimize lost sales. Because of its nature, this policy serves as a benchmark to show the effects of not properly considering the timing uncertainty of demand and the role of inventory timing in managing the latter. After solving all instances, we showed how varying each parameter influences the naïve and optimal inventory strategies. For the optimal policy, we also showed whether a parameter has a stronger influence on the inventory quantity or inventory timing decision. This helps decision makers by clearly identifying which decision must be prioritized depending on the characteristics of the inventory problem they face. We additionally showed, for each parameter, which values lead to the smallest and largest difference in expected profits between the two inventory policies considered, to understand the conditions under which a naïve policy is especially harmful and should be avoided.

Our results showed that the effect on inventory timing is stronger than on inventory quantity for all parameters defining the properties of the stochastic selling season, such as the mean length of the season and the coefficients of variation of the length and start of season, whereas the contrary holds true for the cost parameters of the product, i.e. its critical fractile and the inventory holding costs. Moreover, we showed that avoiding the pitfalls of the naïve policy is especially important when the season's start is highly uncertain, the inventory holding costs are high and the critical fractile is low. These results clearly emphasize the importance of considering both types of demand uncertainties and both inventory decisions simultaneously in the examined setting, in conflict with the lack of studies on the topic. However, several simplifying assumptions were made in the decision model used for the analysis that limit its suitability to solve the inventory problem in real-world settings. The subsequent chapters of the thesis focused on overcoming this limitation.

As opposed to the stylized setting studied in Chapter 3, in real-world settings companies often produce multiple products, each of which could have a selling season with different characteristics, which additionally share the limited capacity of a production resource. Moreover, in many industries, before the production of a product can start, a setup operation which costs money and time must be performed. In addition, in the time horizon which the firm needs to consider in its production/inventory planning problem, the products might have multiple seasons, and each of them is likely to not have a clear-cut pattern. Therefore, in Chapter 4, we examined how the inventory problem in these more complex but practically relevant settings can be tackled.

Consequently, we parted from the continuous time model used in Chapter 3, and modeled the problem using a discrete time stochastic capacitated lot-sizing problem (SCLSP). The objective considered also shifted from profit maximization to cost minimization subject to service-level constraints, a more common approach to inventory planning in practice. This widely used model allowed us to capture the aforementioned more complex production setting and to consider more complex demand patterns. At the same time, modeling the problem in discrete time and defining the demand distribution for each product-period combination requires to correctly capture the auto- and cross-correlations of demand, which are naturally present in a setting with seasonal demand as the one analyzed. Indeed, the shape of the season clearly connects the demand of different periods for a given product, and the demand of different products is correlated if they share the same selling season or have different but interconnected seasons.

We showed that a popular approach to solve the SCLSP found in the literature consists of solving the scenario approximation of the problem. In the analyzed setting, using this solution method requires to build scenarios of demand with the correct auto- and cross-correlations. However, we observed that studies on stochastic lot sizing did not consider settings with seasonal demand, and that most studies further assumed that there were no demand correlations. Furthermore, they did not propose efficient ways to estimate these correlations or to generate scenarios obeying these correlations. Therefore, we proposed an alternative formulation of the scenario approximation of the SCLSP which uses what we define as a cumulative demand scenario (CDS) approach. In contrast to the standard approach found in the literature that defines and uses scenarios which are paths of demand in a scenario tree, the CDS approach requires the generation of scenarios of cumulative demand.

We first compared these two approaches theoretically, to show that they lead to an equivalent estimation of the stochastic variables of the problem and highlighted the advantages of the newly proposed approach. These consist of disregarding demand dependencies when generating scenarios and, therefore, not needing to estimate these complex dependencies in the first place. This is clearly useful in practice, because, first, not many scenario-generation methods are able to generate scenarios of dependent demands without a considerable increase in the complexity of the scenario-generation process, and, second, with the limited amount and imperfect quality of data available in many practical settings, the estimation of dependencies is a challenging task. We then compared the performance of the classic and CDS approaches in a numerical study under the common assumption in the literature of dynamic and uncorrelated demand. First, the results showed that our approach is computationally efficient, because it can obtain high-quality solutions even with a limited number of scenarios. Second, as opposed to the classic approach using common sampling techniques, the results showed that the quality of the solutions is robust to changes in all the parameters of the problem, most notably the coefficient of variation of demand, even with a limited number of scenarios. We finally discussed the limitations of the approach, which, however, we believe to not have a significant impact on its potential benefits. These benefits should be especially large when the CDS approach is applied in a setting with seasonal and uncertain demand – that is the topic of the last chapter of this thesis.

In Chapter 5, we applied the CDS approach in a real-world case study to perform an inventory analysis for an agrochemical company. We started by describing the problem setting analyzed, which is the typical one of the agrochemical industry already introduced in Chapter 2. We chose the stochastic counterpart of the well-known and flexible general lot-sizing and scheduling problem (GLSP) to model the inventory problem of the company, motivated by the complexity of the production processes and the presence of uncertainty. Because of the seasonal and uncertain nature of demand for the products, we used the CDS approach, presented in Chapter 4, to approximate and solve the problem.

We then discussed several additional modeling choices in accordance with the needs and goals of the company, before presenting the model formulation of the problem. Out of all the parameters of the problem, the most challenging ones to obtain were the demand scenarios, because of the limited amount of past data. This lack of sufficient data necessary to estimate a well-defined demand distribution for each product is inherent in the market of CPPs, due to the yearly seasonality of the demand and the issue that past data are not representative of the demand in future seasons, because of the changing portfolio of the market players and of the different customers. Therefore, we discussed the method used to overcome this issue and generate scenarios of cumulative demand starting from past demand and forecast data. This involved using past cumulative relative errors and pooling them over products, assuming that they are independently and identically distributed (iid) random variables.

Finally, we presented the results of the analysis. We showed that the high uncertainty of demand, combined with the high service-level requirements of the industry, necessitates a large investment in safety stock, which makes this inventory driver more important than seasonality and economies of scale. This result emphasizes that it is especially valuable for the company to devote resources and time to improve both its planning strategy under uncertainty and the accuracy of its demand forecasts. Moreover, it shows that even a small increase in the target level can cause a substantial increase in the safety stock needed to reach that target, emphasizing the importance of a careful choice of the target service level.

Furthermore, we demonstrated how our model can facilitate clearly visualizing the trade-off between inventory costs and service level by applying and simulating the performance of the developed model in past seasons. In addition, we used past seasons' data to show how the model compares to the planning method currently used by the company. Given that our model outperformed the existing method, we subsequently analyzed the difference between the plans of the two methods in detail and clearly identified the sources of this difference. Our results showed that the major drawback to the company's current approach is that it lacks a systematic and statistically-based procedure to determine the level of safety stock necessary to reach its target service level. The other important drawback is that it is a sequential approach which tries to deal with the uncertainty of demand in a first step by determining safety stocks, and then develops a production plan to meet these inventory targets in a second step – these two steps are carried out independently from one another, thus neglecting that the optimal safety stock is a function of the production plan. Therefore, the company could benefit by addressing either of the two pitfalls of their planning procedure, and our model is a promising solution to address both.

Because we assumed in our analysis that the relative cumulative forecast errors of all products are iid, we later showed that the CDS approach is robust to misspecifications of the demand distribution by using a different distribution assumption. Moreover, we investigated the effects of changing the target service level definition from an aggregate level to a product-specific level. We showed that this change leads to an only relatively small increase in safety stocks, and that the optimal distribution of service levels among products using aggregate service-level constraints is not excessively uneven. In the last step of the analysis, we proposed a profit-maximizing version of the SGLSP-CDS model developed. We simulated and compared the performance of these two models in past seasons and showed that the profit-maximizing model can support the company in choosing the service-level targets of individual products.

In this thesis, we made multiple valuable contributions to the literature on the inventory management problem of a firm selling products with a seasonal and uncertain demand, both in terms of quantity and timing. However, the presented results are not without limitations and combined with a general shortage of studies on this specific, challenging and highly relevant topic, they provide several future research opportunities.

In the analysis of Chapter 3, the interpretation of the effect of certain problem's parameters on the optimal inventory decisions remains open. Future studies could improve the quality and accuracy of the results obtained by solving the scenario approximation of the problem using a grid-search solution approach. Also, additional practically relevant parameter settings could be considered. Finally, assumptions that limit the applicability of the model to solve real-world problems could be relaxed. Most notably, the model could be extended to consider multiple products sharing a single production resource with limited capacity, in order to investigate what parameters of the products and their seasons' characteristics would influence their priority in terms of safety stock investment and production sequence. Although we move in this direction in the remaining two chapters of the thesis, continuing to use a stylized model would help isolate the effects of changing the problem's parameters to develop a deeper and richer theoretical understanding of the problem and the trade-offs therein, because it provides control and transparency unlike more complex models.

The CDS approach was shown to be computationally efficient and robust to changes in the parameters of the problem. However, future studies could examine the exact reasons for these valuable properties of this approach. Future studies could also conduct a comparison, similar to that made in Chapter 4, between the CDS approach and the standard approach that is used in the literature, where a different scenario-generation method is applied in the standard approach that can account for demand dependencies, in a setting with auto- and cross-correlations. Moreover, the performance of the CDS approach in a rolling horizon (RH) planning framework is also worth investigating, because of the widespread use of this planning strategy in practice.

Concerning the real-world case study presented in Chapter 5, different directions for future research which could provide further valuable insights to the company, can be identified. Extending the model to consider other steps in the supply chain process of the company, such as the production of active ingredients and the bottling and labeling of formulated products, would be an interesting research direction to pursue. In addition, the performance of the CDS approach could be tested by applying the planning model under a RH strategy, in order to provide a better estimate of its performance. Therefore, this would also be a valid option to consider if the model is applied in actual future planning efforts. Moving away from the CDS approach, both for evaluation and future planning purposes, the model could be extended to a multi-stage one to consider re-planning opportunities within the planning horizon and to obtain better solutions than the single-stage SGLSP-CDS model proposed. However, this would lead to both an increase in complexity of the model and the need to estimate more complex demand distributions than those needed by the CDS approach.

To maintain the advantage of the CDS approach, a valid opportunity to still consider re-planning opportunities, even though less accurately, is to extend the CDS approach by following the methodology presented by Tavaghof-Gigloo and Minner (2020). This involves applying the single-stage model in a RH fashion, taking into account future re-planning opportunities by appropriately decreasing the safety stock needed to achieve the given target service level. As another alternative to the CDS approach, the problem could be modeled using the standard scenario-approximation methodology of stochastic programming coupled with scenario-generation techniques which can create scenarios as paths through a scenario tree. This would allow to control risk by measuring the variance of the cost objective function, which is a useful tool in practice. However, this would necessitate estimating the auto- and cross-correlations of demand and generating scenarios with the correct dependencies, a difficult task in this real-world setting, which the CDS approach was specifically designed to avoid.

In conclusion, we trust that the work conducted in this thesis will serve as a stepping stone for further research into the challenging and relevant inventory management problem faced by firms serving stochastic selling seasons.

A Numerical test details

Problem instances have 5, 10 or 20 products, as well as 5, 10 or 20 time periods. In order to obtain the dynamic demand time series, the procedure followed by Helber et al. (2013) was the following. First for each product a different average demand $E[D_k]$ was specified. Second, to obtain a dynamic demand, for each period t the expected demand of product k, $E[D_k t]$, was defined by drawing a random number from a normal distribution with mean $E[D_k]$ and standard deviation $VC^{ip} \cdot E[D_k]$, where VC^{ip} is the coefficient of variation of the expected demand. Two time series of expected demand were defined: one setting $VC^{ip} = 0.2$, creating a moderately dynamic series, and another setting $VC^{ip} = 0.3$, creating a more volatile series. We used the two time series of expected demand obtained, following this procedure, by Helber et al. (2013). To define demand uncertainty, the demand of product k in period t is assumed to be normally distributed with mean $E[D_{kt}]$ and (time-invariant) standard deviation $\sigma_{kt} = \sigma_k = \overline{E[D_k]} \cdot VC^d$, where $\overline{E[D_k]} = \frac{\sum_T E[D_{kt}]}{T}$ and VC^d is the coefficient of variation of demand. In the test instances this coefficient is set equal to 0.1 or 0.3, to simulate a higher and lower demand forecasting accuracy, respectively. To obtain test instances with K products and T periods, the first K products and the first T periods are considered, respectively. The holding costs per unit h_k and the processing time per unit pt_k are set equal to 1, and overtime costs oc equal to 100 per unit of overtime. The setup costs for product k are defined as $sc_k = \frac{\overline{E[D_k]} \cdot TBO^2 \cdot h_k}{2}$, where $TBO = \{1, 2, 4\}$ is the time-between-orders parameter. The setup time of product k, st_k, is defined as $st_k = ts^{rel} \cdot \overline{E[D_k]} \cdot pt_k$, where $ts^{rel} = \{0, 0.25\}$ is the setup time as a fraction of the period processing time. The capacity parameter, b_t is dynamic and defined as $b_t = \frac{\sum_{k \in \mathcal{K}} pt_k \cdot E[D_{kt}]}{Util}$, where $Util = \{0.6, 0.75\}$ is the utilization due to processing. The δ service-level target is varied across instances: it is set equal to 0.8, 0.9 or 0.95 for all products in a given instance. Finally, the sample size S used by all sampling methods is set equal to 10, 30 or 50.

B Additional analysis results

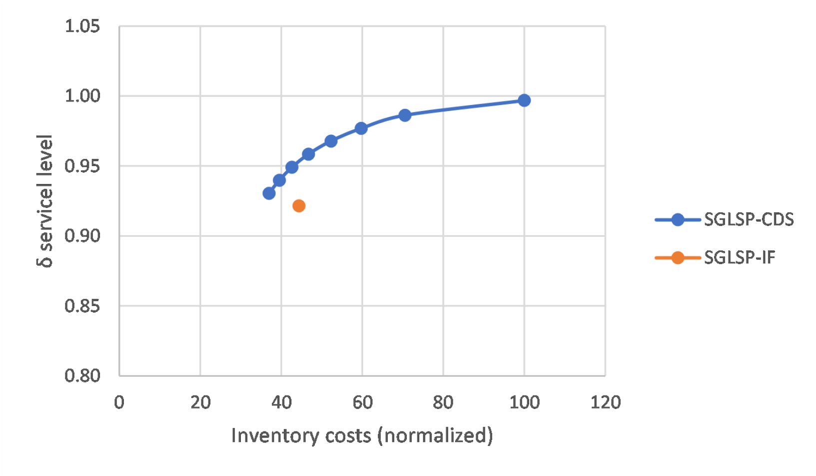

Figure 42: Expected performance of the SGLSP-CDS and GLSP-IF models- Season 2

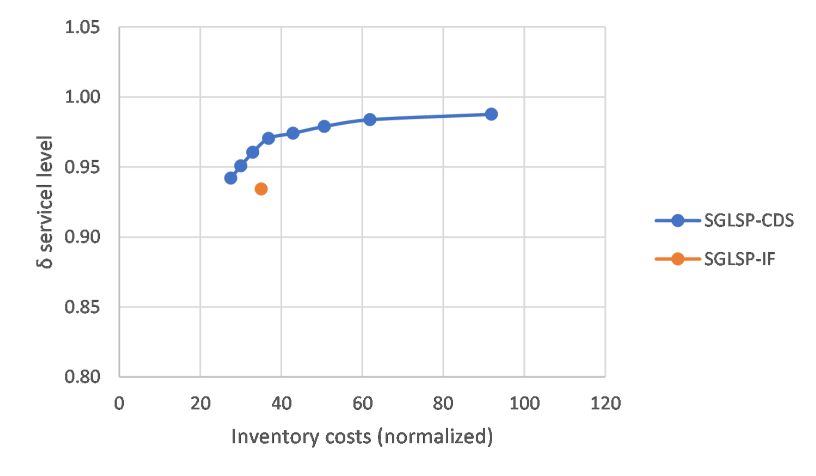

Figure 43: Actual performance of the SGLSP-CDS and GLSP-IF models – Season 2

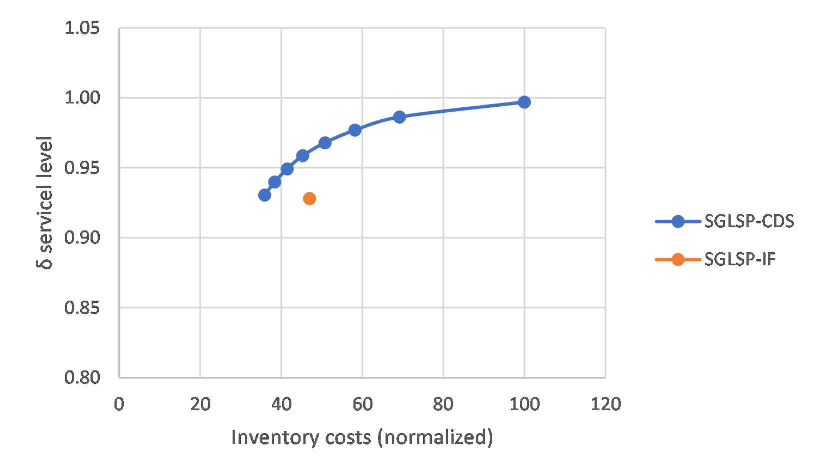

Figure 44: Expected performance of the SGLSP-CDS and GLSP-IF models – Season 3

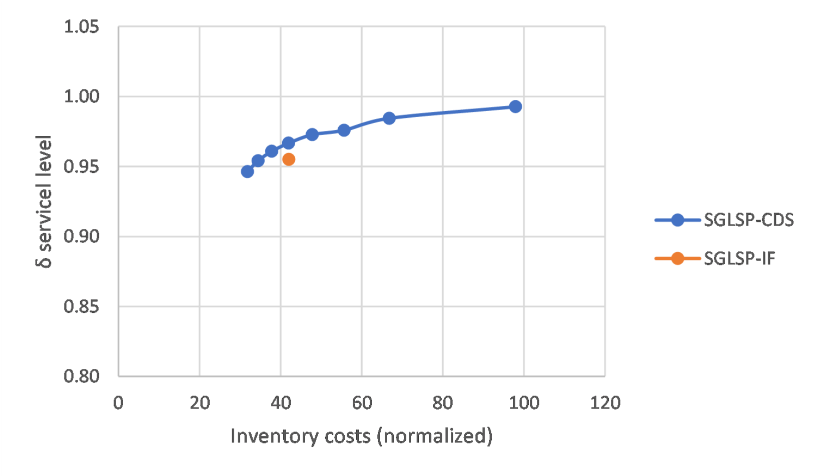

Figure 45: Actual performance of the SGLSP-CDS and GLSP-IF models – Season 3

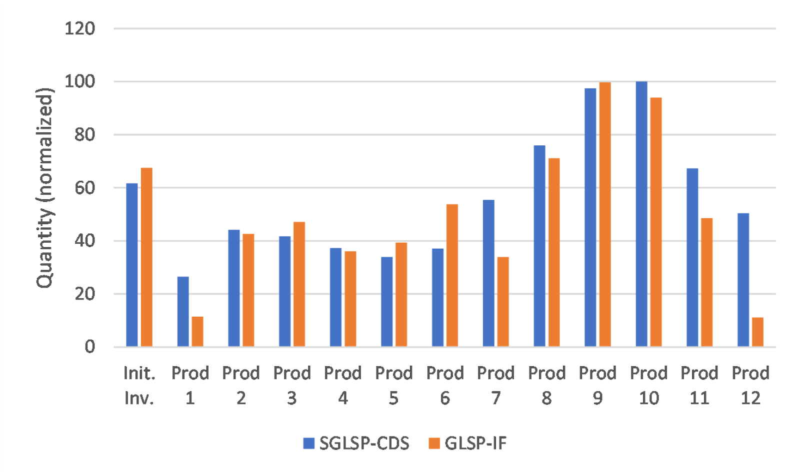

Figure 46: Comparison of optimal plans between SGLSP-CDS and GLSP-IF models – Season 2

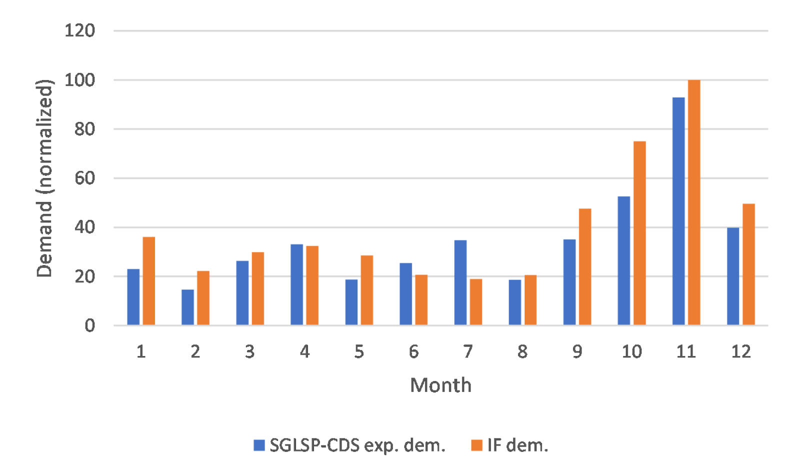

Figure 47: Comparison of input demand between SGLSP-CDS and GLSP-IF models – Season 2

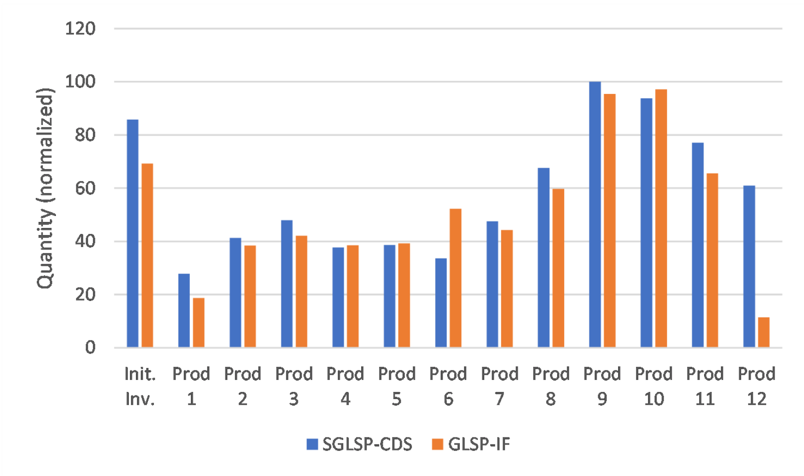

Figure 48: Comparison of optimal plans between SGLSP-CDS and GLSP-IF models – Season 3

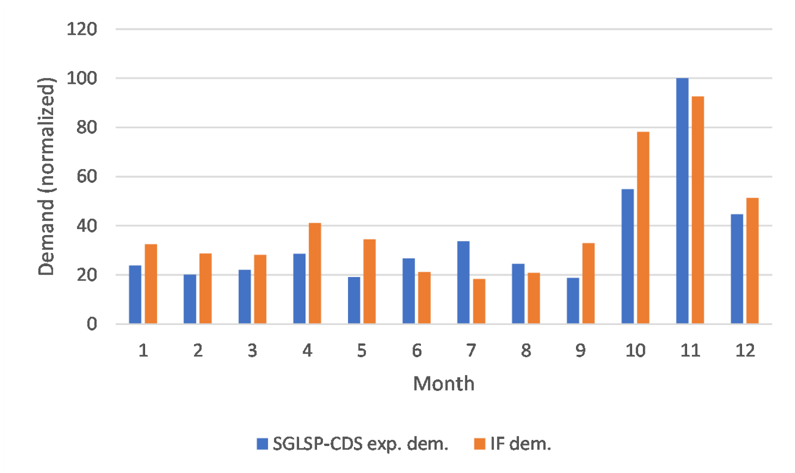

Figure 49: Comparison of input demand between SGLSP-CDS and GLSP-IF models – Season 3

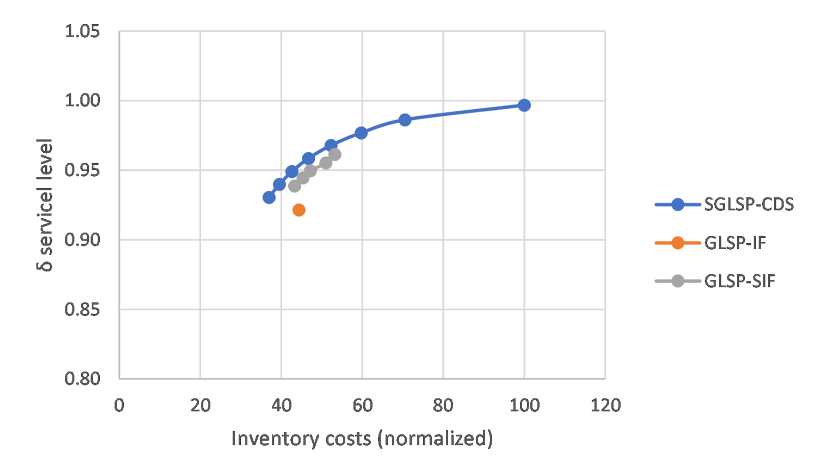

Figure 50: Comparison of SGLSP-CDS, GLSP-IF and GLSP-SIF models – expected performance in Season 2

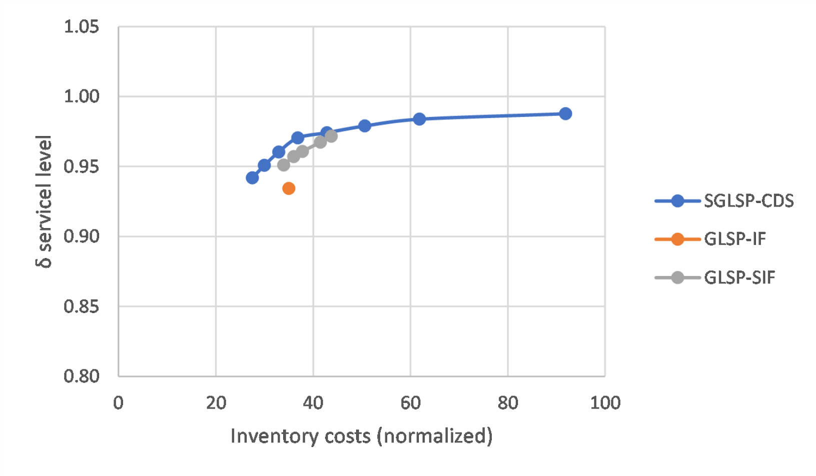

Figure 51: Comparison of SGLSP-CDS, GLSP-IF and GLSP-SIF models – actual performance in Season 2

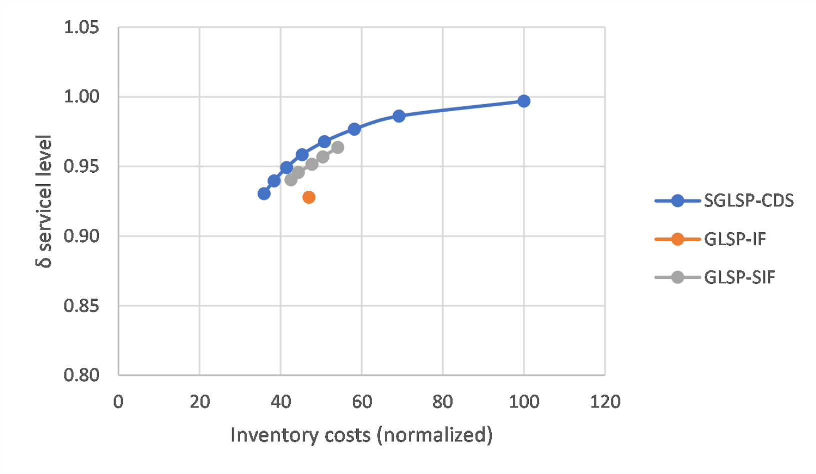

Figure 52: Comparison of SGLSP-CDS, GLSP-IF and GLSP-SIF models – expected performance in Season 3

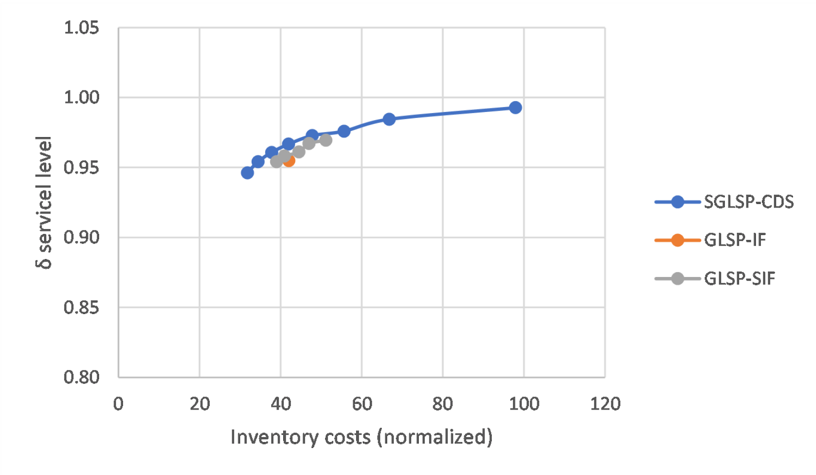

Figure 53: Comparison of SGLSP-CDS, GLSP-IF and GLSP-SIF models – actual performance in Season 3

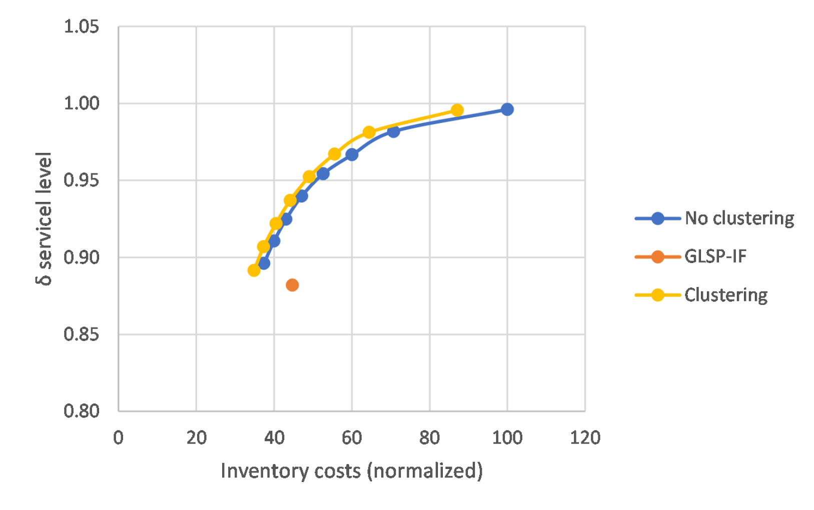

Figure 54: Robustness check to misspecifications of the probability distribution of demand – Season 2

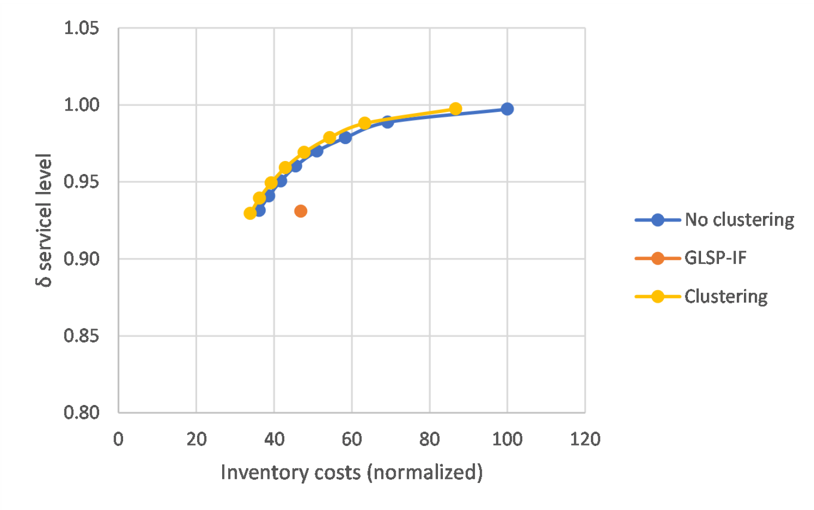

Figure 55: Robustness check to misspecifications of the probability distribution of demand – Season 3

C Mathematical formulation of the inventory problem of Phase 3

Phase3 Model

$$\text{Min} \sum_{k\in\mathcal{K}}\sum_{t\in\mathcal{T}} h_k \cdot IP_{kt} + \sum_{k\in\mathcal{K}}\sum_{n\in\mathcal{N}}\sum_{m\in\mathcal{M}} h_k \cdot x_{knm} \tag{111}$$

Subject to

$$\sum_{n\in\mathcal{N}_t}\sum_{k\in\mathcal{K}} pt_{km}\cdot x_{knm} + \sum_{n\in\mathcal{N}_t}\sum_{k\in\mathcal{K}}\sum_{j\in\mathcal{K}} st_{kjm}\cdot z_{kjnm} \le b_{mt} \qquad \forall m\in\mathcal{M},\ \forall t\in\mathcal{T} \tag{112}$$

$$x_{knm} \le \frac{b_{mt}}{pt_{km}}\cdot y_{knm} \qquad \forall k\in\mathcal{K},\ \forall m\in\mathcal{M},\ \forall n\in\mathcal{N} \tag{113}$$

$$x_{knm} \ge l_{km}\cdot(y_{knm} - y_{k,n-1,m}) \qquad \forall k\in\mathcal{K},\ \forall m\in\mathcal{M},\ \forall n\in\mathcal{N} \tag{114}$$

$$\sum_{k\in\mathcal{K}} y_{knm} = 1 \qquad \forall m\in\mathcal{M},\ \forall n\in\mathcal{N} \tag{115}$$

$$z_{kjnm} \ge y_{k,n-1,m} + y_{jnm} - 1 \qquad \forall k,j\in\mathcal{K},\ \forall m\in\mathcal{M},\ \forall n\in\mathcal{N} \tag{116}$$

$$IP_{kt} = IP_{k0} + \sum_{n\in\mathcal{N}_T}\sum_{m\in\mathcal{M}} x_{knm} - CD_{kt} \qquad \forall k\in\mathcal{K},\ t=1 \tag{117}$$

$$IP_{kt} = IP_{k0} + \sum_{n\in\mathcal{N}_T}\sum_{m\in\mathcal{M}} x_{knm} + \sum_{\tau=1}^{t-1}\sum_{n\in N_\tau}\sum_{m\in\mathcal{M}} x_{knm} - CD_{kt} \qquad \forall k\in\mathcal{K},\ t\ge 2 \tag{118}$$

$$\sum_{n\in\mathcal{N}}\sum_{m\in\mathcal{M}} x_{knm} \ge CD_{kT} \qquad \forall k\in\mathcal{K} \tag{119}$$

$$x_{knm} \ge 0 \qquad \forall k\in\mathcal{K},\ \forall m\in\mathcal{M},\ \forall n\in\mathcal{N} \tag{120}$$

$$y_{knm} \in \{0,1\} \qquad \forall k\in\mathcal{K},\ \forall m\in\mathcal{M},\ \forall n\in\mathcal{N} \tag{121}$$

$$z_{kjnm} \ge 0 \qquad \forall k,j\in\mathcal{K},\ \forall m\in\mathcal{M},\ \forall n\in\mathcal{N} \tag{122}$$

$$IP_{kt} \ge 0 \qquad \forall k\in\mathcal{K},\ \forall t\in\mathcal{T} \tag{123}$$

Bibliography

Albrecht, M., J. Rohde, and M. Wagner (2015). Master planning. In H. Stadtler, C. Kilger, and H. Meyr (Eds.), *Supply chain management and advanced planning: concepts, models, software, and case studies* (5 ed.)., pp. 155–176. Springer.

Alem, D., E. Curcio, P. Amorim, and B. Almada-Lobo (2018). A computational study of the general lot-sizing and scheduling model under demand uncertainty via robust and stochastic approaches. *Computers and Operations Research 90*, 125–141.

Aloulou, M. A., A. Dolgui, and M. Y. Kovalyov (2014). A bibliography of non-deterministic lot-sizing models. *International Journal of Production Research 52*, 2293–2310.

Arrow, K. J., T. Harris, and J. Marschak (1955). Optimal inventory policy. *Econometrica 19*, 250–272.

Azoury, K. and J. Miyaoka (2013). Managing production and distribution for supply chains in the processed food industry. *Production and Operations Management 22*(5), 1250–1268.

Bassett, M. (2000). Assigning projects to optimize the utilization of employees' time and expertise. *Computers & Chemical Engineering 24*, 1013–1021.

Bassett, M. (2018). Optimizing the design of new and existing supply chains at Dow AgroSciences. *Computers and Chemical Engineering 114*, 191–200.

Bassett, M. and L. Gardner (2013). Designing optimal global supply chains at Dow AgroSciences. *Annals of Operations Research 203*, 187–216.

Birge, J. R. and F. Louveaux (2011). *Introduction to stochastic programming* (2nd ed.). Springer.

Bloemen, R. and J. Maes (1992). A DSS for optimizing the aggregate production planning at Monsanto Antwerp. *European Journal of Operational Research 61*, 30–40.

Böcker, T. G. and R. Finger (2017). A meta-analysis on the elasticity of demand for pesticides. *Journal of Agricultural Economics 68*, 518–533.

Bookbinder, J. and J. Y. Tan (1988). Strategies for the probabilistic lot-sizing problem with service-level constraints. *Management Science 34*, 1096–1108.

Boulaskil, Y., J. C. Fransoo, and E. N. G. van Halm (2009). Setting safety stocks in multi-stage inventory systems under rolling horizon mathematical programming models. *OR Spectrum 31*, 121–140.

Bouma, E. (2003). GEWIS, a weather-based decision support system for timing the application of plant protection products. *EPPO Bulletin 33*, 483–487.

Brandimarte, P. (2006). Multi-item capacitated lot-sizing with demand uncertainty. *International Journal of Production Research 4*(15), 2997–3022.

Buschkühl, L., F. Sahling, S. Helber, and H. Tempelmeier (2010). Dynamic capacitated lot-sizing problems: A classification and review of solution approaches. *OR Spectrum 32*(2), 231–261.

Cachon, G. and C. Terwiesch (2012). *Matching supply with demand: An introduction to Operations Management* (3 ed.). McGraw Hill.

Calfa, B., A. Agarwal, I. Grossmann, and J. Wassick (2014). Data-driven multi-stage scenario tree generation via statistical property and distribution matching. *Computers and Chemical Engineering 68*, 7–23.

Cario, M. C. and B. L. Nelson (1997). Modeling and generating random vectors with arbitrary marginal distributions and correlation matrix. Technical report, Department of Industrial Engineering and Management Science, Northwestern University, Evanston, (IL, USA).

Chen, S. and D. Su (2022). A multi-stage stochastic programming model of lot-sizing and scheduling problems with machine eligibilities and sequence-dependent setups. *Annals of Operations Research 311*, 35–50.

Choi, T.-M., D. Li, and H. Yan (2004). Optimal single ordering policy with multiple delivery modes and bayesian information updates. *Computers and Operations Research 31*, 1965–1984.

Conejo, A., M. Carrion, and J. Morales (2010). *Decision Making Under Uncertainty in Electricity Markets*. Springer.

Copil, K., M. Wörbelauer, H. Meyr, and H. Tempelmeier (2017). Simultaneous lotsizing and scheduling problems: A classification and review of models. *OR Spectrum 39*, 1–64.

Curcio, E., P. Amorim, Q. Zhang, and B. Almada-Lobo (2018). Adaptation and approximate strategies for solving the lot-sizing and scheduling problem under multistage demand uncertainty. *International Journal of Production Economics 202*, 81–96.

De Smet, N., S. Minner, E. Aghezzaf, and B. Desmet (2020). A linearisation approach to the stochastic dynamic capacitated lotsizing problem with sequence-dependent changeovers. *International Journal of Production Research 58*(16), 4980–5005.

Deler, B. and B. L. Nelson (2000). Modeling and generating multivariate time series with arbitrary marginals using an autoregressive technique. Technical report, Department of Industrial Engineering and Management Science, Northwestern University, Evanston, (IL, USA).

Desai, G. M. (1970). Factors determining demand for pesticides. *Economic and Political Weekly 5*, A181–A183.

Dessouky, M., B. Kijowski, and S. Verma (1999). Simultaneous batching and scheduling for chemical processing with earliness and tardiness penalties. *Production and Operations Management 8*, 433–444.

Di Domenica, N., C. Lucas, G. Mitra, and P. Valente (2009). Scenario generation for stochastic programming and simulation: A modelling perspective. *IMA Journal of Management Mathematics 20*, 1–38.

Drexl, A. and K. Haase (1995). Proportional lotsizing and scheduling. *International Journal of Production Economics 40*, 73–87.

Elimam, A. A. (1995). A decision support system (DSS) for agricultural pesticide production planning. *European Journal of Operational Research 81*, 17–34.

Eppen, G. D. and R. K. Martin (1987). Solving multi-item capacitated lot-sizing problems using variable redefinition. *Operations Research 35*(6), 832–848.

European Commission (2016, Jun). Frequently asked questions: endocrine disruptors. https://ec.europa.eu/commission/presscorner/detail/en/MEMO_16_2151.

European Commission (2020, Nov). Pesticides. https://ec.europa.eu/food/plant/pesticides_en.

Fisher, M. and A. Raman (1996). Reducing the cost of demand uncertainty through accurate response to early sales. *Operations Research 44*(1), 87–99.

Fleischmann, B. (1990). The discrete lot-sizing and scheduling problem. *European Journal of Operational Research 44*, 337–348.

Fleischmann, B. and H. Meyr (1997). The general lotsizing and scheduling problem. *OR Spectrum 19*, 11–21.

Fleischmann, B., H. Meyr, and M. Wagner (2015). Advanced planning. In H. Stadtler, C. Kilger, and H. Meyr (Eds.), *Supply chain management and advanced planning: Concepts, models, software, and case studies* (5 ed.)., pp. 71–95. Springer.

Fritz, M. and T. Hausen (2009). Electronic supply network coordination in agrifood networks: Barriers, potentials, and path dependencies. *International Journal of Production Economics 121*, 441–453.

Groebner, D. F. and C. M. Merz (1990). Solving the inventory problem for the sale of seasonal merchandise. *Journal of Small Business Management 28*, 19–26.

Gupta, Y., P. Sundararaghavan, and M. U. Ahmed (2003). Ordering policies for items with seasonal demand. *International Journal of Physical Distribution & Logistics Management 33*, 500–518.

Haase, K. (1996). Capacitated lot-sizing with sequence dependent setup costs. *OR Spectrum 18*(1), 51–59.

Hadley, G. (1962). Generalizations of the optimal final inventory model. *Management Science 8*, 454–457.

Harris, F. W. (1913). How many parts to make at once. *Factory, The Magazine of Management 10*(2), 135–6.

Helber, S., F. Sahling, and K. Schimmelpfeng (2013). Dynamic capacitated lot sizing with random demand and dynamic safety stock. *OR Spectrum 35*(1), 75–105.

Hilger, T., F. Sahling, and H. Tempelmeier (2016). Capacitated dynamic production and remanufacturing planning under demand and return uncertainty. *OR Spectrum 38*, 849–876.

Hoyland, K. and S. W. Wallace (2001). Generating scenario trees for multistage decision problems. *Management Science 4*(2), 295–307.

Hu, Z. and G. Hu (2016). A two-stage stochastic programming model for lot-sizing and scheduling under uncertainty. *International Journal of Production Economics 180*, 198–207.

Hu, Z. and G. Hu (2018). A multi-stage stochastic programming for lot-sizing and scheduling under demand uncertainty. *Computers and Industrial Engineering 119*, 157–166.

Karmarkar, U. S. and L. Schrage (1985). The deterministic dynamic product cycling problem. *Operations Research 33*, 326–345.

Kaut, W. (2014). A copula-based heuristic for scenario generation. *Computer Management Science 11*, 503–516.

Kaut, W. and S. Wallace (2003). Evaluation of scenario-generation methods for stochastic programming. *Pacific Journal of Optimization 3*(2), 257–271.

Kaut, W. and S. Wallace (2011). Shape-based scenario generation using copulas. *Computer Management Science 8*, 181–199.

Koca, E., H. Yaman, and M. S. Aktürk (2015). Stochastic lot sizing problem with controllable processing times. *Omega 53*, 1–10.

Lang, J. C. (2010). *Production and inventory management with substitutions*. Springer.

Lau, H.-S. and A. H.-L. Lau (1997). Reordering strategies for a newsboy-type product. *European Journal of Operational Research 103*, 557–572.

Li, L., S. Song, C. Wu, and R. Wang (2017). Fix-and-optimize and variable neighborhood search approaches for stochastic multi-item capacitated lot-sizing problems. *Mathematical Problems in Engineering 2017*.

Li, Y. and G. Hu (2017). Shop floor lot-sizing and scheduling with a two-stage stochastic programming model considering uncertain demand and workforce efficiency. *Computers and Industrial Engineering 111*, 263–271.

Liu, S. and L. G. Papageorgiou (2013). Multiobjective optimisation of production, distribution and capacity planning of global supply chains in the process industry. *Omega 41*, 369–382.

Liu, S. and L. G. Papageorgiou (2018). Fair profit distribution in multi-echelon supply chains via transfer prices. *Omega 80*, 77–94.

Loretan, M. (1997). Generating market risk scenarios using principal components analysis: Methodological and practical considerations. In K. Arrow, S. Karlin, and P. Suppes (Eds.), *The Measurement of Aggregate Market Risk, CGFS Publications No.7*, pp. 23–61. Bank for International Settlements.

Maravelias, C. T. and I. E. Grossmann (2001). Simultaneous planning for new product development and batch manufacturing facilities. *Industrial Engineering & Chemical Research 40*, 6147–6164.

McGraw, K. E. and M. M. Dessouky (2001). Sequence-dependent batch chemical scheduling with earliness and tardiness penalties. *International Journal of Production Research 39*, 3085–3107.

Meistering, M. and H. Stadtler (2017). Stabilized-cycle strategy for capacitated lot sizing with multiple products: Fill-rate constraints in rolling schedules. *Production and Operations Management 26*, 2247–2265.

Mitchell, W. C. (1927). *Business cycles: the problem and its setting*. NBER. https://www.nber.org/books-and-chapters/business-cycles-problem-and-its-setting.

Mitra, G. and N. Di Domenica (2010). A review of scenario generation methods. *Int. J. Computing Science and Mathematics 3*(3), 226–244.

Montgomery, D. C. (2017). *Design and analysis of experiments* (9 ed.). Wiley.

Mooney, P. (2018). Blocking the chain. http://dx.doi.org/10.17169/refubium-2748.

Nahmias, S. (1977). On ordering perishable inventory when both demand and lifetime are random. *Management Science 24*, 82–90.

Nishimoto, R. (2019). Global trends in the crop protection industry. *Journal of Pesticide Science 44*, 141–147.

Phillips, M. W. A. (2020). Agrochemical industry development, trends in R&D and the impact of regulation. *Pest Management Science 76*, 3348–3356.

Ravindran, A. (1972). Management of seasonal style-goods inventories. *Operations Research 20*, 265–275.

Rogers, J. (1958). A computational approach to the economic lot scheduling problem. *Management Science 4*(3), 264–291.

Rosenzweig, C., A. Iglesias, X. B. Yang, P. R. Epstein, and E. Chivian (2001). Climate change and extreme weather events - implications for food production, plant diseases, and pests. *Global Change and Human Health 2*, 90–104.

Rossi, R., O. A. Kilic, and S. A. Tarim (2015). Piecewise linear approximations for the static-dynamic uncertainty strategy in stochastic lot-sizing. *Omega 50*, 126–140.

Sahling, F., L. Buschkühl, H. Tempelmeier, and S. Helber (2009). Solving a multi-level capacitated lot sizing problem with multi-period setup carry-over via a fix-and-optimize heuristic. *Computers and Operations Research 36*(9), 2546–2553.

Saliby, E. (1990). Descriptive sampling: A better approach to Monte Carlo simulation. *The Journal of the Operational Research Society 41*(12), 1133–1142.

Scarf, H. (1959). The optimality of (S, s) policies in the dynamic inventory problem. In K. Arrow, S. Karlin, and P. Suppes (Eds.), *Mathematical methods in the social sciences*, pp. 196–202. Stanford Univeristy Press.

Schlapp, J. and M. Fleischmann (2020). Inventory timing: How to serve a stochastic season. Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3747789.

Schmidt, C. W. and I. E. Grossmann (1996). Optimization models for the scheduling of testing tasks in new product development. *Industrial & Engineering Chemistry Research 35*, 3498–3510.

Schnelle, K. D. (2000). Preliminary design and scheduling of a batch agrochemical plant. *Computers and Chemical Engineering 24*, 1535–1541.

Schütz, P., A. Tomasgard, and S. Ahmed (2009). Supply chain design under uncertainty using sample average approximation and dual decomposition. *European Journal of Operational Research 199*, 409–419.

Schweitzer, M. E. and G. P. Cachon (2000). Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Science 46*, 404–420.

Shah, N. (2004). Pharmaceutical supply chains: Key issues and strategies for optimisation. *Computers and Chemical Engineering 28*, 929–941.

Shapiro, A. (2003). Monte Carlo sampling methods. In A. Ruszczyński and A. Shapiro (Eds.), *Stochastic Programming*, Volume 10 of *Handbooks in Operations Research and Management Science*, pp. 313–344. Elsevier.

Silver, E. A., D. F. Pyke, and D. J. Thomas (2016). *Inventory and Production Management in Supply Chains* (4th ed.). Taylor & Francis Group.

Sousa, R., N. Shah, and L. G. Papageorgiou (2008). Supply chain design and multilevel planning - an industrial case. *Computers and Chemical Engineering 32*, 2643–2663.

Sparks, T. C. and B. A. Lorsbach (2017). Perspectives on the agrochemical industry and agrochemical discovery. *Pest Management Science 73*, 672–677.

Subrahmanyam, S., J. F. Pekny, and G. V. Reklaitis (1994). Design of batch chemical plants under market uncertainty. *Industrial Engineering & Chemical Research 33*, 2688–2701.

Subramanian, D., J. F. Pekny, and G. V. Reklaitis (2000). A simulation-optimization framework for addressing combinatorial and stochastic aspects of an r&d pipeline management problem. *Computers & Chemical Engineering 24*, 1005–1011.

Taş, D., M. Gendreau, O. Jabali, and R. Jans (2018). A capacitated lot sizing problem with stochastic setup times and overtime. *European Journal of Operations Research 273*, 146–159.

Tavaghof-Gigloo, D. (2019). *Essays on dynamic production and safety stock planning.* Dissertation, Technische Universität München, München.

Tavaghof-Gigloo, D. and S. Minner (2020). Planning approaches for stochastic capacitated lot-sizing with service level constraints. *International Journal of Production Research 59*, 5087–5107.

Tempelmeier, H. (2011). A column generation heuristic for dynamic capacitated lot sizing with random demand under a fill rate constraint. *Omega 39*(6), 627–633.

Tempelmeier, H. (2013). Stochastic lot sizing problems. In J. M. Smith and B. Tan (Eds.), *Handbook of Stochastic Models and Analysis of Manufacturing System Operations*, pp. 313–344. Springer New York.

Tempelmeier, H. and S. Herpers (2010). ABC β - a heuristic for dynamic capacitated lot sizing with random demand under a fill rate constraint. *International Journal of Production Research 48*(17), 5181–5193.

Tempelmeier, H. and T. Hilger (2015). Dynamic capacitated lot sizing with random demand and dynamic safety stock. *OR Spectrum 59*, 119–125.

Tempelmeier, H., M. Kirste, and T. Hilger (2018). Linear programming models for a stochastic dynamic capacitated lot sizing problem. *Computers and Operations Research 91*, 258–259.

Wagner, H. and T. M. Whitin (1958). Dynamic version of the economic lot size model. *Management Science 5*, 89–96.

Wang, T., A. Atasu, and M. Kurtuluş (2012). A multiordering newsvendor model with dynamic forecast evolution. *Manufacturing & service operations management 14*, 472–484.

Wang, T. and B. Tomlin (2009). To wait or not to wait: Optimal ordering under lead time uncertainty and forecast updating. *Naval Research Logistics 56*, 766–779.

Zheng, M., Y. Shu, and K. Wu (2015). On optimal emergency orders with updated demand forecast and limited supply. *International Journal of Production Research 53*, 3692–3719.

Curriculum Vitae

Luca Biscaro

02/2017-09/2021	**Doctoral Candidate** University of Mannheim, Mannheim, Germany Chair of Logistics and Supply Chain Management
09/2013-03/2016	**Mannheim Master in Management (M. Sc.)** University of Mannheim, Mannheim, Germany
10/2010-07/2013	**International Business (B. Sc.)** Cracow University of Economics, Cracow, Poland