Aalto University
School of Science
Degree Programme in Engineering Physics and Mathematics

Turkka Anttonen

Procurement Contracts under Uncertain Demand

Master's Thesis
Espoo, October 13, 2016

Supervisor:	Professor Ahti Salo
Advisor:	M.Sc. (Tech.) Susanna Alaja

Aalto University
School of Science
Degree Programme in Engineering Physics and Mathematics

ABSTRACT OF MASTER'S THESIS

Author:	Turkka Anttonen
Title:	Procurement Contracts under Uncertain Demand
Date:	October 13, 2016 — **Pages:** vi + 45
Major:	Systems and Operations Research — **Code:** F3008
Supervisor:	Professor Ahti Salo
Advisor:	M.Sc. (Tech.) Susanna Alaja

At present, high-tech manufacturers experience short product life cycles, high demand uncertainty, and diminishing profit margins. This presents a problem to the procurement department. It has to ensure supply of components but overstocking consumes low profit margins by tying capital in inventories. In order to be able to purchase components at a lower price, the manufacturer may take a bigger share of the risk imposed by demand uncertainty by making a purchasing commitment.

In this thesis, demand uncertainty is modelled by using the Martingale Method of Forecast Evolution on real-life data of a manufacturer in consumer electronics industry. A procurement contract with a quantity commitment, in exchange for a discount, is then compared to a reference contract with a fixed unit price. Monte Carlo simulation is used to derive demand scenarios and procurement-related cost is calculated for each contract. Comparison of the costs provides a tool for assessing whether the manufacturer should enter a commitment contract or not.

The decision greatly depends on the parameters of the contract and the capability to affect them by negotiating. While the method provides multiple indicators that summarize the data, there are also visualizations that help decision makers understand the distributions of possible outcomes.

Keywords:	Procurement contracts, demand uncertainty, cost modeling, Monte Carlo method, visualization
Language:	English

Aalto-yliopisto
Perustieteiden korkeakoulu
Teknillisen fysiikan ja matematiikan koulutusohjelma

DIPLOMITYÖN TIIVISTELMÄ

Tekijä:	Turkka Anttonen		
Työn nimi: Hankintasopimukset ja kysynnän epävarmuus			
Päiväys:	13. lokakuuta 2016	**Sivumäärä:**	vi + 45
Pääaine:	Systeemi- ja operaatiotutkimus	**Koodi:**	F3008
Valvoja:	Professori Ahti Salo		
Ohjaaja:	Diplomi-insinööri Susanna Alaja		

Teknologiateollisuudessa tuotteiden elinkaaret ovat nykyään lyhyitä, kysyntä epävarmaa ja voittomarginaalit pieniä. Tämä muodostaa haasteen yrityksen hankintaosastolle. Sen tulee varmistaa komponenttien saatavuus kasvattamatta kuitenkaan varastoja liiallisesti, jotta sitoutunut pääoma ei heikennä voittomarginaalia entisestään. Pystyäkseen hankkimaan komponentit edullisemmin, valmistava yritys voi sitoutua hankkimaan määrätyn määrän komponentteja ja kantaa näin suuremman osan kysynnän epävarmuuteen liittyvästä riskistä.

Tässä työssä kysynnän epävarmuutta mallinnetaan soveltamalla kulutta-jaelektroniikan valmistajan kysyntädataan martingaalimenetelmää (Martingale Method of Forecast Evolution). Hankintasopimusta, jossa ostaja saa alennuksen sitoutuessaan ostamaan määrätyn määrän komponentteja, verrataan kiinteän yksikköhinnan sopimukseen. Kysyntäskenaariot luodaan Monte Carlo -simulaatiolla ja hankintaan liittyvät kustannukset lasketaan kullekin sopimukselle. Näiden kustannusten vertailu tukee ostajan päätöksentekoa, kun sitoumuksen tekoa harkitaan.

Päätössuositus riippuu sopimuksen parametreista ja siitä, pystyykö ostaja vaikuttamaan niihin neuvottelemalla. Menetelmän avulla saadaan tuotettua sekä tilastollisia jakaumia ja riskiä kuvaavia tunnuslukuja että mahdollisten lopputulemien jakaumia havainnollistavia kuvaajia.

Asiasanat:	Hankintasopimukset, kysynnän epävarmuus, kustannusmallinnus, Monte Carlo -menetelmä, visualisointi
Kieli:	Englanti

Acknowledgements

I would like to express my gratitude to my instructor Susanna Alaja and all my former colleagues at Nokia and Microsoft. More than anything, this thesis will be a memory of the priceless experience of working with and learning from a group of extremely talented people.

I would also like to thank professor Ahti Salo for supervising the thesis. I firmly believe that Aalto University, and especially Systems Analysis Laboratory, did give me the best starting point possible for a professional career and I am truly grateful for that.

Finally, I would like to thank my family for all the support throughout the years. I simply could not have done this without you.

Espoo, October 13, 2016

Turkka Anttonen

Contents

Chapter 1

Introduction

1.1 Background

In high-tech electronics industry, there has been recent development towards shorter product life cycles and more volatile customer demands. In the highly competitive market, demand can decrease dramatically within a short period of time for instance when a competitor introduces a competitive product into the market. On the other hand, if a product turns out to be more successful than expected, catching the upside opportunities may be extremely difficult as some of the components have very long lead times compared to the overall product life cycle. Accurate planning can hence be said to have an important role in the success of a company.

Great demand uncertainty has become well known to some companies. In 2001, Cisco wrote down $2.2 billion mostly in raw materials and components due to weakened demand (Burrows 2003). In late 2011, Research In Motion (RIM) wrote down $485 million as a result of worse than expected demand for their PlayBook products (Research In Motion 2011). Some companies have started adopting new tools and models to address the fact that demand is uncertain. This happens typically through incorporating uncertainty into sales and operations planning (S&OP) by presenting demand uncertainty using different demand scenarios, e.g. pessimistic, most likely, and optimistic

(Sodhi and Tang 2010). If certain demand distribution is assumed and supply chain related cost and revenue parameters defined, planning can be supported by solving an expected profit maximization problem where excess inventory buildup and unmet demand yield penalties.

For a manufacturing company that designs products and purchases materials for assembly, customer demand uncertainty is a big concern because it is very likely to carry risk related to excess inventory and unmet demand. An end product requires all components in its bill-of-material to be present at the time of assembly or otherwise the product can't be assembled. Lacking one component may thus lead to both unmet demand for the end product and excess inventory for other components already ordered and allocated to the production. On the other hand, overstocking can lead into carrying excess inventory all the time consuming the profits. At the end of the product life cycle, there is also the risk of having obsolete inventory which can't be used for anything and thus must be scrapped.

"The [procurement] department of a company is responsible for obtaining the materials, supplies, and services needed to produce a product or provide a service." (Stevenson 2009) Purchasing price is one of procurement departments' main concerns. It is in most cases the biggest cost element when total procurement cost is broken down. Purchased parts typically represent 40 to 60 per cent of an end product's value (Ballou 2004). Especially for companies in high technology industry, material cost can represent a high share of overall cost of goods sold. Through the leverage effect of purchasing, it is easy to justify such focus – a percentage-wise small decrease in the purchasing cost can lead into a considerable increase in profit margin (Wisner 2011). However, many companies have realized that lowest purchasing price does not always mean lowest total cost. It may be useful to consider total landed cost which not only considers the purchasing price but also the transportation cost, warehousing costs, etc. (Erhun and Tayur 2003). The modelling can be taken even further with methods such as total cost of ownership which also includes quality, technology, and support related costs into the model

(Ellram 1995).

Purchasing price is typically agreed as a part of the supply contract which states the terms of business between two contracting parties. The supply contract also defines other conditions than the purchasing price. In order to reduce the risk caused by downstream demand uncertainty, the two parties can e.g. agree that a certain minimum quantity is purchased. In return, the supplier may give a discount on the purchasing price compensating the manufacturer for bearing more risk. This kind of a contract is called a quantity commitment contract.

The decision to enter a quantity commitment contract is made before the realized demand can be observed. The buyer has a forecast of the future demand and can assess the financial sensibility of the commitment based on it. However, if the forecast does not explicitly take into account the uncertainty related to the demand, the buyer may make a commitment which in the light of the most recent forecast may seem beneficial but is actually expected to yield a worse result than an alternative lower-risk contract due to uncertain demand. This is further made worse by the fact that both undercommitting and overcommitting will lead into additional cost.

1.2 Objectives

The objective of this thesis is to develop a framework for a buyer for assessing different supply contracts' financial sensibility under uncertain demand. This includes both quantifying and visualizing the impact of demand uncertainty in an otherwise deterministic procurement problem. Different ways of modelling the demand uncertainty are reviewed as well as different kinds of supply contracts that are commonly used.

The method is applied into a real-life scenario using a method based on the Martingale Method of Forecast Evolution to generate different demand scenarios. Two supply contracts, a price only and a quantity commitment contract, are studied by modelling the cost incurred in both using deterministic cost functions that take demand scenarios as an input. The decision whether to enter the commitment contract is a single-stage decision. Monte Carlo simulation is used to create the scenarios, calculate costs, and aggregate results.

1.3 Structure

Chapter 2 reviews the literature on demand uncertainty, supply contracting, and procurement cost modelling. Related mathematical models are also examined. The model used in the thesis is presented in chapter 3 together with the mathematical model and methods. Chapter 4 contains the results of the analysis. Summary and conclusions are in chapter 5.

Chapter 2

Literature review

2.1 Demand uncertainty

Demand uncertainty is one of the forms of uncertainty that any company faces continuously. Customers either do not order or are able to cancel their preliminary (soft) orders until very late in the demand fulfillment process. Demand uncertainty causes problems for the operations trying to make demand and supply meet. In industries where products are standard and have a steady demand, finished goods can be produced to stock (make-to-stock) to provide buffer against demand fluctuations. At the other end of the spectrum, in industries where products are highly customized for customers and thus have a higher demand volatility, products cannot be made to stock because capital may be tied in finished goods for a long time. In this situation, products are typically made to order. When operating in make-to-order mode, the manufacturer often needs to plan production capacity prior to learning the realized demand. The manufacturer may keep some flexibility in the production capacity and have inventories of raw materials and semi-finished goods in order to be prepared for demand fluctuations. However, unused capacity and excessive safety stocks incur cost that suppy chain management aims at minimizing.

Uncertain demand is especially difficult for products with a short life cycle. When the life cycle is short, there is increasing pressure to fulfill all demand without having inventory left at the end of the life cycle. For example in fashion industry, markdowns can be used to increase demand at the end of the selling season at the cost of profit. In the worst case, the manufacturer or retailar may be left with obsolete goods which have to be scrapped at a further cost. End-of-life management deals with consuming inventories so that no obsolete stock is left. With short life cycle products, the problem is that a decrease in demand can lead to such a sudden end-of-life that only managing the supply chain reactively is not sufficient to avoid accumulating obsolete stock.

A retail environment facing a similar problem has lent its name to this type of a problem in supply chain management literature. The newsvendor problem is motivated by a newsvendor who needs to order a certain amount of newspapers prior to the selling season. In case of a daily newspaper, the selling season is assumed to be one day as the availability of the next day's newspaper and news would obsolete its value. Running into a stockout and failing to meet the demand means that the newsvendor has lost profit. In addition to the basic dynamics, lots of extensions have been formulated in the newsvendor framework making the literature in this area vast and diverse (see e.g. Porteus 1990).

Demand uncertainty can be modelled in many ways. Typically, an assumption on the distribution is made and the parameters of that distribution are estimated from data. Normal distribution due to its ease is used commonly. Granularity of demand representation varies between different models. In some analyses, demand for a product may be considered as the lifetime demand and little attention is paid on how the demand is distributed within the lifetime. Operationally and tactically, it is crucial to understand when the demand will realize, i.e. whether majority of the demand occur at the beginning of or whether it is stable over the whole lifetime. With both objectives, it is also relevant to define what the time horizon in question is. The planning horizon is typically finite in a real-life setting but in an ideal situation an infinite horizon may also be considered. A rolling horizon is also a common type of a horizon used in real life. At the beginning of period 1, rolling horizon of t periods considers periods $[1, t]$. Decisions are made for first period and after observing the outcome, horizon and the problem is revised for periods $[2, t+1]$ with the new information from period 1. Chand et al. (2002) have compiled a literary review of different horizons' use in different problems in operations management literature.

Different ways of modelling demand and related uncertainty include the use of time series, assuming stationary demand, assuming a forecast update model, using probabilities to model a soft order turning into a hard one in forecast sharing setting, or creating scenarios. A common way to consider stochastic demand is to assume independent and identically distributed demand like the (truncated) normal distributions in Bassok and Anupindi (1997). Hausman (1969) studied real life forecast updates and concluded that in many of them the ratio of successive forecasts conformed to the lognormal distribution. Graves et al. (1986) follows a similar setup but with the absolute forecast update being the random variable. The individual forecast updates for different months were assumed independent. Heath and Jackson (1994) present the model as Martingale Model of Forecast Evolution (MMFE) considering both the additive and multplicative models with lognormal distribution. In this model, forecast update rounds were assumed independent but interdependence was allowed within a single update. Kaminsky's (2004) model is based on updating a band that narrows down over time to a more accurate demand forecast. The model assumes that the level of realized demand exists within the band. Contrary to this assumption, Cattani and Hausman (2000) argue that forecasts do not always improve over time. Terwiesch et al. (2005) uses a logit model to model probablities of soft orders turning into firm orders in a forecast sharing setting. Sodhi (2005) uses a binary tree built on an autoregressive time series model of the first order, AR(1), to create demand scenarios. The binary tree includes 2^T scenarios where T is the number of time periods in the time horizon. Nagali et al. (2008) mention that

Hewlett-Packard use regression analysis of historical forecasts and shipments to model demand uncertainty. Gaur et al. (2007) empirically show that the dispersion among experts' forecasts provides a good measure for demand volatility. They obtained the result by studying the standard deviation of forecast error so this only suggests that experts' forecasts can be used to understand the magnitude of uncertainty while they don't necessarily improve forecast accuracy or impact the underlying process.

Companies can manage demand uncertainty with different approaches. Gupta and Maranas (2003) distinguish between a shaper and an adopter strategy. In the former, the company attempts to change the distribution of the uncertain demand and thus minimize the downside risk while retaining the upside potential. In the adopter strategy, the company tries to control the risk exposure rather than influence the demand. Literature is skewed towards the adopter strategy side. Fisher and Raman (1996) describes a method of using early sales data to enhance the accuracy of demand forecast in fashion industry. Wu et al. (2006) study leading indicators in semiconductor industry in order to gain more accurate forecasts. Nagali et al. (2008) report that Hewlett-Packard use a portfolio of contracts with their suppliers to hedge against demand fluctuations.

Supply chain risk management deals with an array of interrelated risks. Uncertain customer demand is only one of these risks. Tang (2006) provides a good review of supply chain risks. Other procurement related risks include uncertainty in cost, yield, lead times, capacity, liquidity, and exchange rates among others.

2.2 Supply contracting

Supply contracts are made to define the rules governing the business relationship between a buyer and a supplier. Anupindi and Bassok (1999) state that supply contracts should capture three types of flows: material, information, and financial. Each of these have different parameters that further describe the agreement between the buyer and the supplier. As Tsay et al. (1999) point out, the purpose of a supply contract is also to have the two parties share the risk related to different uncertainties of a supply chain. By agreeing on the quantity to be supplied and the price paid for it, the buyer may mitigate a risk related to price fluctuations and uncertain capacity. On the other hand, the supplier is guaranteed capacity utilization and a selling price.

Much of operations research on supply contracting is focused on finding a contract that allows supply chain (or channel) coordination. Supply chain coordination means that either a party or the contractual form and parameters incentivise the parties to act in a way that is optimal for the whole supply chain. A typical cause of inefficiency in a supply chain is asymmetric information e.g. in the form of buyer's private information about her demand. A coordinating contract might be one that e.g. disincentivizes inflating shared forecasts beyond likely figures in order to make the supplier build more capacity that could be later used to catch unforeseen upsides.

Supply contracts take various forms stating how the risk and revenue are shared between the buyer and the supplier. Some of the most common types of supply contracts are presented below.

> The most simple form of a supply contract is a **price only** or **fixed price** contract. The buyer and supplier agree on a wholesale price. Once the transaction is completed, the buyer owns the material but is not further obliged to anything. (Lariviere 1999)
>
> **Fixed price with incentives** is an extension for the fixed price contract that rewards the supplier for work that exceeds the agreed standard. It guarantees the supplier a minimum price but also incentivizes her to strive for exceptional performance. (Van Weele 2009)
>
> In a **cost-plus** contract, the buyer pays the supplier her cost plus a percentage guaranteeing a profit to the supplier. This is

effectively risk-free for the supplier while the buyer carries more risk than in a fixed price contract. (Van Weele 2009)

In a **quantity discount** contract, the supplier offers the buyer a lower per unit purchase price for purchasing more units. The quantity discount can be an **all-units** or an **incremental quantity discount**. In the former, the price of all purchased units is reduced when a certain quantity threshold is reached. In the latter, the discount is only applied to the units exceeding the threshold. (Burnetas, Gilbert, and Smith 2007)

A **buy back** contract is a two-stage system. In the first stage, the buyer purchases the material from the supplier at a wholesale price. If she is left with excess inventory after she has satisfied all her demand, the supplier buys back the excess stock at a reduced price. (Lariviere 1999)

The buyer in a **revenue sharing** contract purchases at a wholesale price but also awards the supplier a percentage of her revenue. (Cachon 2003)

Total minimum quantity commitment contract refers to a contract in which the buyer commits to purchase a certain minimum quantity during a fixed period of time. Usually the supplier offers a price discount for such commitment and is able to do so because her risk is reduced due to the commitment. If there is a possibility to purchase more than the initial commitment quantity, the contract is called **total minimum quantity commitment with flexibility**. The flexibility beyond the commitment quantity may be available only at a premium. (Anupindi and Bassok 1999)

Total minimum dollar volume commitment is typically used if the commitment is made for multiple products. Also in this case, the supplier offers a price discount in return but the commitment may change within the products included in the contract terms as long as their combined dollar value equals at least the agreed amount. (Anupindi and Bassok 1999)

Periodical commitments take many forms. E.g. a **rolling horizon flexibility** contract has the buyer committing to certain periodical quantities at the beginning of the horizon. The commitment is then revised periodically within certain limits of flexibility for the future periods. Another form is **periodical commitment with options**. Similar to the rolling horizon flexibility contract, initial commitments are made but the buyer also purchases options that she can later decide to exercise. (Anupindi and Bassok 1999)

Forecast sharing is closely related to supply contracting. Forecast sharing means that the buyer shares her forecast for the forecasting horizon with the supplier. This information will help the supplier plan her production and make decisions on capacity investments. Typically, the shared forecast is not a binding commitment to purchase. It is known that the buyer has an incentive to inflate the forecast she communicates to the supplier in order to get the supplier to build more capacity for her, thus reducing risk related to not having enough capacity (Cachon and Lariviere 2001). However, knowing this the supplier will not build enough capacity. In a long-term relationship, there may be an incentive to share the forecast truthfully (Ren, Cohen, Ho, and Terwiesch 2010).

In addition to contracting, many studies include an alternative source for goods, the spot market. The spot market is open market available to all buyers where prices and availability are defined by the market dynamics. Hence, there is considerable uncertainty in both price and supply. The spot market only exists for commodity goods such as electricity, or standardized electronic components like memory chips. Nevertheless, it may be a prominent part of the purchasing strategy especially if the prices have high volatility because it requires no commitments and purchases can be made at any time provided that there is supply. Martínez-de Albéniz and Simchi-Levi (2005) have included the spot market into their supply model.

2.3 Procurement cost modelling

Procurement cost modelling is a topic that all relevant research has to take a stand on. Depending on the focus of the research, different elements of the overall supply chain cost may be included. It is often the case that procurement cost is modelled as a part of a company's profit maximization problem. In the simplest form, one can only take into account the purchasing cost that the buyer has to pay in return for the goods it purchases. However, this is a very narrow view of the cost. Waters (2003) lists reorder cost, holding cost and shortage cost as other cost elements to consider. Reorder cost refers to the cost related to ordering, following up on the order, physically receiving and inspecting the goods. Holding cost consists of the cost of capital tied in the stock, storage cost, stock loss and obsolescence, handling, administration and insurance. Holding cost is estimated to be approximately 25% of the value of the held item annually. Shortage cost is considered when there is a stockout that prevents from satisfying customer demand. It may come in the form of re-scheduling, paying a premium for rushed deliveries, or using more expensive suppliers. All of these cost items may be seen as procurement-related costs if the supply contract is such that it may incur holding or shortage cost. In addition to these, salvage value is also used in some models. It is the value that the item has e.g. on the after-market when the manufacturer no longer has use for it. It can also be negative e.g. in case the left-over stock can't be sold and has to be scrapped at additional cost. Another type of cost is what Callioni et al. (2005) call component devaluation cost. Especially in high-tech industries where product life cycles are short, the introduction of replacing technology may result in dramatic drops in the prices of the replaced technology. In their example, a central processing unit (CPU) experienced up to 40% price reduction during its nine-month life cycle.

Different models include different cost elements into the cost function that is typically the objective function for the optimization problem. Bassok and Anupindi (1997) only consider purchasing price, holding and storage cost. Gupta and Maranas (2003) do not limit their problem to the procurement cost so they also have some production related costs. Procurement related are material unit cost and inventory holding costs as well as penalties for safety stock violations and unmet demand. In most cases, the cost elements included in the model are purchasing, holding, and penalty cost (Anupindi and Bassok 1999; Bassok and Anupindi 2008; Cachon and Zipkin 1999; Chen and Krass 2001). Penalty cost mainly refers to backlogging (backordering) penalty. Backlogging of orders, i.e. whether customer orders that could not be fulfilled will be carried over to the next period, is not present in all models. If no backlogging is assumed, such orders can be interpreted as automatically cancelled and will incur penalty cost. Other kinds of penalties and their parameters, defined in the supply contract, are also subject to research when supply chain coordination is sought after (Cachon and Zipkin 1999).

2.4 Mathematical methods

A variety of mathematical methods have been used in supply chain and procurement related problems. Incorporating stochasticity into the model usually implies the use of stochastic programming if the problem is formulated as an optimization problem.

Bassok and Anupindi (1997) use dynamic recursion to reduce a quantity commitment problem to a standard newsvendor problem for which the solution is well known. Further, they conduct a computational study assuming normal distribution. They run a sensitivity analysis by changing the parameters of the model and solving the problem repeatedly.

Heath and Jackson (1994) model forecast evolution by fitting a MMFE model into forecast data. The multiplicative MMFE is based on an observation that in a forecast updating scheme, the ratios of successive forecasts follow a lognormal distribution. Successive forecasts of a certain demand in the future are assumed to form a martingale process, i.e. the expected value of each updated forecast is the same as the previous one. Using past data, the variance-covariance matrix is estimated. The setting is a multi-product multi-period setting where demand is assumed neither stationary nor independently distributed but dependencies are allowed e.g. between different products' forecasts. Using the variance-covariance matrix and initial state vector (matrix), future forecasts can be generated. Heath and Jackson used the methodology in a safety stock optimization problem but the linear programming model is a separate component of the overall simulation model so MMFE as a method may also be used in a different setting.

Hausman and Peterson (1972) also make the assumption of lognormally distributed ratios of successive forecasts. Their problem is minimizing the costs due to overproduction and underproduction. They formulate a dynamic program but in the absence of an analytical solution defer to studying the performance of heuristic strategies.

Chen and Krass (2001) study an optimal commitment quantity. They analytically define the strategy in their setting, but note that with additional cost elements, the problem would be a lot harder to solve. They also present a numerical example in which they do not use optimization methods.

Supply chain models often experience a problem emerging from the complexity of the model combined with potentially multiple sources of uncertainty. Hence the analytical solutions are difficult, or even impossible, to obtain (Shapiro 2001). In these cases, simulation is an option that can be considered. While a simulation does not yield the optimal solution in analytical form, it can still provide valuable information and in some cases be the only way to optimize. It is also usually easy to create a simulation model and to modify it according to the specific scenario, e.g. the supply contract in question. Out of the stochastic simulation methods, Monte Carlo method is probably best known. Wu and Olson (2008) use Monte Carlo simulation in their optimization problem of multi-objective supplier selection. First, 600 scenarios are created using the distributions assumed for the random variables, then the optimization algorithm is run for all of them and at the end, the results are aggregated. Van Landeghem and Vanmaele (2002) also use Monte Carlo simulation in supply chain context. The problem is a multi-echelon inventory problem but as they note, the tactical-level supply chain modelling easily becomes very complex and analytical solutions may be very impractical to obtain.

Chapter 3

Modelling the problem and data

3.1 Supply chain

The supply chain used in the procurement cost modelling problem of this thesis is made up of two or three parties: a manufacturer and one or two suppliers. The manufacturer purchases components from the suppliers and assembles them into products that it subsequently sells to its customers. The products are high-tech products with a relatively short life cycle. Due to the high level of customer-specific customization, the products are made to order. The customer demand that the manufacturer faces is uncertain. A monthly-level forecast for a rolling 12-month horizon is revised at the beginning of each month by planning experts who use past data as well as quantitive and qualitative information about the future demand to do so. At the end of the month, the manufacturer learns the realized demand but has already fulfilled all demand that she is able to. This means that the component inventory held at the beginning of the month t, $i(t)$, the quantity purchased during month t, $q(t)$, and realized demand during month t, $d(t)$, must satisfy

$$i(t) + q(t) \geq d(t) - U(t), \tag{3.1}$$

where $U(t)$ denotes the unmet demand during period t. As it is assumed that left-over inventory is carried to the next period and demand is not backlogged, the conditions can be written as

$$i(t+1) = \max\left\{i(t) + q(t) - d(t), 0\right\} \tag{3.2}$$

$$U(t) = \max\left\{d(t) - i(t) - q(t), 0\right\}. \tag{3.3}$$

When the end of the last month where $d(t) > 0$ is reached, the product life cycle has ended from production point-of-view. The salvaged inventory $s(t)$ is thus defined as

$$s(t) = i(t), \text{when } d(t) = 0. \tag{3.4}$$

It can be assumed that $d(t) = 0$ implies $d(T) = 0 \forall T > t$. This means that the manufacturer does not need to carry the obsolete inventory but can salvage it immediately after the production has ended. However, it is assumed that the inventory can't be salvaged until $d(t) = 0$ even if the cumulative demand forecast does not exceed current inventory when $d(t) > 0$.

The manufacturer and both suppliers, A and B, are fully flexible and have unlimited capacity. However, the suppliers do have a higher cost level due to having such flexibility. If the manufacturer can commit to a certain quantity, the suppliers can reduce their cost level by selling the extra capacity to other customers or by optimizing their production scheduling. As a result, the capacity will be constrained as defined in the supply contract. The components that the manufacturer buys from these suppliers are custom-designed and cannot be manufactured by other companies so it is fair to assume that a spot market does not exist. An annual devaluation of 10% is assumed.

There are no information or material flow lead times in the supply chain. The focus is on tactical level decisions so also other operational concerns are disregarded, such as lot-sizing, ordering, production scheduling, yield, etc.

3.2 Model selection

As discussed in section 2.1, there are many ways of modelling uncertain demand. Out of these methods, many, including e.g. Bassok and Anupindi (1997), are not suitable because of the non-stationary nature of the demand. The method of Heath and Jackson (1994) seems most appropriate. It can be applied to a similar forecast updating model that the manufacturer uses. It also allows for dependencies between different products' demands. This is a desirable property because the products can be seen as competing against each other or reinforcing the demand of other products through the common brand. Both positive and negative correlations are thus possible.

The goal of this thesis is to develop a framework and methods for assessing supply contracts from the procurement point-of-view. Thus, the objective is to find the contract parameters that minimize procurement related cost. In many other applications in the same field, the objective is maximizing profit.

Procurement cost function is formulated based on the parameters and form of the supply contracts in question. The cost function is deterministic while the only stochastic element is demand.

The demand model and deterministic procurement cost function are connected using Monte Carlo method of simulation. Thus, it is possible to use a deterministic cost function without modeling it as a stochastic one. This approach is also modular, which is advantageous from practical applications' point-of-view. Either the demand scenario creation method or cost function may be changed into a different one without having to change the other. Such a scenario asset may also be used outside of the procurement context, e.g. in factory capacity or logistics network simulations.

3.3 Forecast updating and demand realization

Forecast updating follows the method in Heath and Jackson (1994). First, the log ratios are calculated

$$v(\tau, n, t) = \log\big(d(\tau, n, t)\big) - \log\big(d(\tau - 1, n, t)\big), \tag{3.5}$$

where $d(\tau, n, t)$ refers to the forecast made at the beginning of month τ for product n's demand in month t. For $\tau > t$, d refers to realized demand. Further, the ratio of successive forecasts, $R(\tau, n, t)$ can be defined and the updating formula for forecasts written as

$$R(\tau, n, t) = e^{v(\tau,n,t)} \tag{3.6}$$

$$d(\tau, n, t) = R(\tau, n, t)d(\tau - 1, n, t). \tag{3.7}$$

Next, an infinite vector $v(\tau)$ is constructed so that the first N components contain $v(\tau, 1, \tau), \ldots, v(\tau, N, \tau)$, the next N are $v(\tau, 1, \tau{+}1), \ldots, v(\tau, N, \tau{+}1)$ and so forth. The infinite vector $v(\tau)$ can be truncated at $12N$ components as 12 is the length of the forecast horizon. The components of $v(\tau)$ are assumed to be jointly normally distributed. However, $v(\tau)$ and $v(\sigma), \tau \neq \sigma$ are independent and identically distributed with a mean of 0 based on the martingale assumption. Now, the variance-covariance matrix $\mathbf{\Sigma}$ can be estimated from the vectors $v(\tau - 1), \ldots, v(\tau - \omega)$ of the past periods.

When updating the forecast for the first time, the number of months in the forecast update vector will match the number of months in the initial plan. However, on the next round, a plan will not exist for the 12th month. In the dataset which is used to derive $\mathbf{\Sigma}$ and the initial plan, the last month's forecast has been made by the planning experts so it does exist. Heath and Jackson did not include such method in their model and it is not done in this thesis. This means that demand $d(\tau, n, t) = 0 \forall \tau \geq \bar{\tau} + 13$ where $\bar{\tau}$ is the last available plan's month.

Also, in case the product's life cycle ends before the end of the forecast horizon, the zero-demands have to be removed as $d(\tau, n, t) = 0$ would require taking $\log 0$ which isn't defined. This is solved by replacing zero-demands with $d(\tau, n, t) = 1$ in the plans. When unmodified demand is greater than zero, it is at minimum in the order of thousands so this modification will not effectively change the total demand.

3.4 Supply contracts

3.4.1 Reference case: price only contract with forecast sharing

The reference supply contract is the kind of contract that the manufacturer uses by default. It is a price only contract, i.e., only the price, $p(t)$, is agreed between the manufacturer and the supplier. The prices are agreed for one quarter at a time at the beginning of the quarter. The manufacturer shares her forecast with the supplier and the supplier manages the inventory at a location close to the manufacturer's production facilities. The manufacturer only purchases the materials against firm customer orders so she does not carry inventory at any point apart from what is operationally necessary to run the operations without line-downs. This stock can be disregarded as insignificant for the total cost.

In this model, the supplier carries practically all of the risk. Hence, it charges a higher unit price for the components, including a risk premium they have seen reasonable. Supplier A offers a considerably lower price and in a normal situation supplier B would only be used in case a disruption occured that would prevent supplier A from supplying. Supplier B is nevertheless able to offer low enough a price so that the product cost of the manufacturer ensures her a profit.

3.4.2 Alternative case: total quantity commitment with flexibility

The alternative case that the manufacturer is assessing is a total quantity commitment with flexibility. Supplier offers a contract where the manufacturer commits to purchase Q units during the next T months. In return, the supplier offers a unit price that is discounted by d compared to the price only contract. The scheduling in those periods is still fully flexible, but if the cumulative purchases at the end of the Tth month have not reached Q units, the difference between the purchased amount and Q is purchased and carried over to the next month as component inventory. When committing to Q units, the supplier plans her own operations accordingly and sets a maximum flexibility of K units that cannot be exceeded during the commitment period. For purchases beyond Q units, the supplier charges an additional premium p_c. It is assumed that the manufacturer first uses this flexibility and only after the full capacity is utilized, uses supplier B or must leave the demand unmet. If $Q+K$ does not enough fulfill the demand during the commitment period, a penalty of p_u is given for the extra units. This can either be viewed as the penalty of unmet demand or the unit price from supplier B.

In this setting, the manufacturer assumes risk related to supply, excess inventory, and cost. Supply-wise, the existence of supplier B means that there is no stockout risk. On the other hand, if supplier B has to be used extensively, there is a risk that the overall purchasing cost will exceed that of the reference case.

3.5 Procurement cost function

In the reference case, the procurement cost C_R is

$$C_R(\tau, n, t) = d(\tau, n, t)p(t), \tag{3.8}$$

where $p(t)$ is the price agreed between the manufacturer and the supplier or the manufacturer's estimate of what the price will be.

In the alternative case, the cost function is formed of five different cost elements. Assuming that such contract has been made, for the whole contract the cost, C_A, is

$$C_A = C_c + C_p + C_u + C_h + C_s, \tag{3.9}$$

where C_c is the purchasing cost for the quantity commitment, C_p is the cost for the overage that can be purchased from supplier A, C_u is the penalty for unmet demand, C_h is the inventory holding cost, and C_s the salvage cost. C_u may also be interpreted as the cost of purchasing from supplier B. Broken further, C_c is simply

$$C_c = (1 - d)p(\bar{\tau})Q, \tag{3.10}$$

where $\bar{\tau}$ is the first month of the commitment period, i.e. the quantity discount for the whole commitment period is assumed to be given relative to the first month's price.

For simplicity, the realized demand during the commitment period is

$$D = \sum_{n \in P} \sum_{j=1}^{T} d(\bar{\tau} + j, n, \bar{\tau} + j - 1), \tag{3.11}$$

where P denotes the list of products which contain the component in question. The demand forecast at the beginning of the commitment period is

$$D' = \sum_{n \in P} \sum_{j=1}^{T} d(\bar{\tau}, n, \bar{\tau} + j - 1). \tag{3.12}$$

Price premium purchases C_p are

$$C_p = (1 + p_c)(1 - d)p(\bar{\tau}) \min \Big\{ \max \big\{ D - Q, 0 \big\}, K \Big\}. \tag{3.13}$$

Components will be purchased at the discounted price until the commitment level Q is reached. Then, up to the maximum capacity $Q + K$, components will be purchased at the discounted price plus upward flexibility premium p_c. If the demand exceeds $Q + K$ units during the commitment period, the exceeding part will be considered unmet demand. The cost function for unmet demand C_u is

$$C_u = p_u \max \big\{ D - K - Q, 0 \big\}, \tag{3.14}$$

where p_u is the penalty paid for unmet demand. As noted in the model description, this penalty can also be interpreted as the fixed unit cost of buying the components from a second, fully flexible supplier B. In that case, p_u is considered low enough for the manufacturer to make a profit. A wide range of values needs to be thus evaluated to conclude the feasibility of the commitment contract in each case as the penalty for unmet demand can be assumed to be considerably higher than the unit price from supplier B.

Inventory holding cost C_h can be written as

$$C_h = \frac{c_h}{12} \sum_{j=0}^{\infty} \left(\frac{i_p(\bar{\tau} + T + j) + i_p(\bar{\tau} + T + j + 1)}{2} p(\bar{\tau} + T + j) \right), \tag{3.15}$$

where $i_p(\bar{\tau} + T + j)$ is the inventory projection made at $\bar{\tau} + T$ for the future month $\bar{\tau} + T + j$. Thus, it is assumed that the holding cost is paid for the mean inventory during the month. The term i_p is defined

$$i_p(\bar{\tau} + T) = \max \left\{ Q - D, 0 \right\} \tag{3.16}$$

$$i_p(\bar{\tau}+T+j) = i_p(\bar{\tau}+T+j-1) - \sum_{n \in P} d(\bar{\tau}+T, n, \bar{\tau}+T+j-1), \text{where } j \geq 1. \tag{3.17}$$

If the purchased quantity exceeds the total left-to-go demand, salvaged inventory cost C_s will be added as in

$$C_s = p_s \max \left\{ i_p(\bar{\tau} + T + j) - \sum_{n \in P} \sum_{j=1}^{\infty} d(\bar{\tau} + T, n, \bar{\tau} + T + j - 1), 0 \right\}. \tag{3.18}$$

The different costs incurred in the three scenarios $D_{1,2,3}$ satisfy

$$D_1 < Q < D_2 < Q + K < D_3 \tag{3.19}$$

are presented in figure 3.1 and the quantities applicable for each cost element are shown in table 3.1 using the earlier notation.

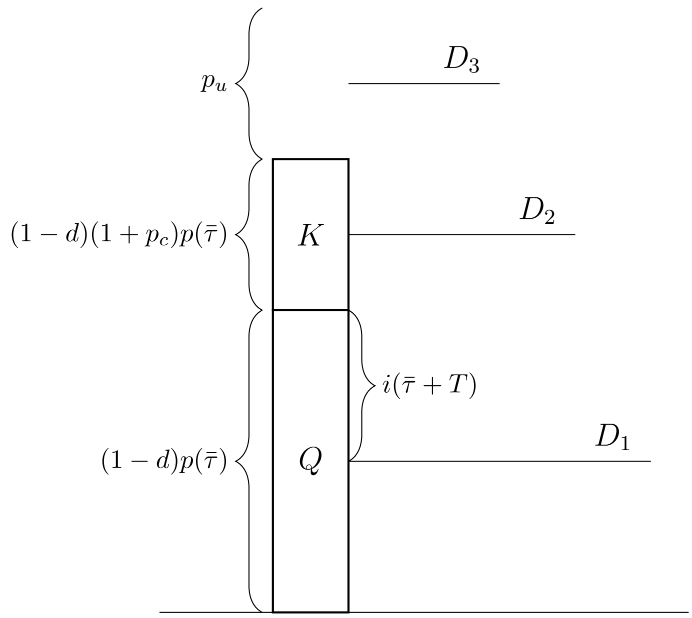

Figure 3.1: Figure showing the commitment Q, flexibility K, prices paid for units purchased on the left, and three demand scenarios described in table 3.1 on the right.

Table 3.1: Quantities for the calculation of different cost elements in three different demand scenarios.

Scenario	D_1	D_2	D_3
Commitment	Q	Q	Q
Flexibility	-	$D_2 - Q$	K
Unmet demand	-	-	$D_3 - K - Q$
Holding	$Q - D_1$ depending on demand after commitment period	-	-
Salvage	$Q - D_1$ depending on demand after commitment period	-	-

3.6 Monte Carlo method

Monte Carlo method is utilized as follows:

1. l_{MC} random samples are generated using the initial states $d(\bar{\tau}, n, t)$ and variance-covariance matrix $\mathbf{\Sigma}$. 12 forecast update vectors are generated for each scenario.

2. Forecast update vectors are applied sequentially following (3.7).

3. Deterministic procurement cost is calculated for the reference case using (3.8) and for the alternative case using (3.9).

4. Results are aggregated.

3.7 Metrics

First, we note that because the procurement related costs can be compared for different quantities, an adjustment has to be made to the reference case's cost function. If the manufacturer commits to a greater quantity than she will consume during the commitment period, she will have some stock left after the commitment period is over. Hence, she does not need to purchase as many units after the commitment whereas in the reference case, this quantity would be eventually (partially) purchased at a price that may be higher than the alternative case's discounted price plus the premium paid for utilizing the flexibility. This is taken into account by adding a term into the cost function (3.8) so that

$$C_{aR} = C_R + \sum_{j=1}^{\infty} \Big((i_p(\bar{\tau} + T + j) - i_p(\bar{\tau} + T + j - 1))p(t + j) \Big). \quad (3.20)$$

The profitability of the alternative case is presented by calculating the mean savings compared to the reference case as a percentage. Savings S and the mean savings $\bar{S}$ are defined as

$$S = 1 - \frac{C_A}{C_{aR}} \quad (3.21)$$

$$\bar{S} = \frac{\sum^{l_{MC}} S}{l_{MC}}, \tag{3.22}$$

where $\sum^{l_{MC}}$ denotes the sum over the scenarios created using the Monte Carlo method. In comparison, the savings generated in the deterministic case where $D = D'$ are also calculated as a reference for the case in which the latest demand forecast is taken for granted at the time of decision making. When $D = D'$, setting $q = 1$ yields $C_A = C_c$ and further $S = d$.

In addition to calculating the mean savings, other metrics are also calculated. Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR) are typical risk metrics used to describe financial risk. VaR is the threshold value that defines the level of return at the point where the cumulative distribution function of the returns equals a certain level, i.e. VaR is a quantile. Levels typically used are 1%, 5% and 10%. VaR is a widely used risk metric but has some undesirable properties from mathematical point-of-view such as not being subadditive (Szegö 2002). CVaR is often used instead of VaR, especially when the objective is to minimize the risk of loss. CVaR is defined as the expected return for the cases that have a lower return than the respective VaR. In the discrete case, these metrics are defined as a quantile of the sample and the mean of the values below that quantile.

To facilitate the comparison of different scenarios, the commitment quantity q is expressed as a percentage of the cumulative demand forecast for the commitment period at the beginning of month $\bar{\tau}$

$$q = \frac{Q}{D'}. \tag{3.23}$$

Because commitment quantity is the value that is easiest to adjust, the results are mainly presented as a function of q. Maximum S^* and the maximum point q^* with respect to $\bar{S}$, i.e.

$$S^* = \max\left\{\bar{S}\right\} \tag{3.24}$$

$$q^* = \arg\max\{\bar{S}\} \tag{3.25}$$

are defined as the expected savings at optimum and the respective commitment quantity. In addition, the mean value $\bar{S}_{100}$ at $q = 100\%$ is calculated representing the expected savings if commitment quantity equals the demand forecast for the commitment period. Also the intervals $[q_l, q_u]$ where $\bar{S} \geq 0$ are defined describing the commitment quantity interval that is expected to yield savings compared to the reference case and the standard deviation of S, σ_S, is calculated. For VaR and CVar, the results are presented as value pairs of maximum points and maxima $q^*_{\text{VaR}}, \text{VaR}^*$ and $q^*_{\text{CVaR}}, \text{CVaR}^*$ respectively.

3.8 Data

The demand data used in this thesis consists of 1872 data points made up from 9 forecast revisions of 12 months' forecasts for 16 products. As described in section 3.3, the realized demands are also included in the data. An example of an evolving forecast horizon is presented in figure 3.2 where the graph is divided into 21 individual plots for the 12-month horizon that is rolled over 9 times. Each block represents a month t and within the blocks, $d(\tau, n, t)$ is plotted as a function of τ. For simplicity, $\tau = 1$ refers to the initial forecast in the dataset whereas $\bar{\tau} = 1$ later on. This example represents an optimistic demand forecast. As new forecasts are generated, in most cases the demand is revised to a lower or equal level compared to the previous forecast.

3.9 Parameters

The parameters are set as follows: the decision point $\bar{\tau}$ is assumed to be the point at which the quarterly price agreement takes place. Purchasing cost $p(\bar{\tau}) = 2$ based on the offer from the supplier. Commitment discount d, flexibility premium p_c, and penalty cost for unmet demand p_u are varied, subject to the constraint $(1 - d)p(\bar{\tau}) \leq (1 - d)(1 + p_c)p(\bar{\tau}) \leq p_u$ so that the manufacturer's strategy of purchasing at the flexibility price, rather than choosing penalty for unmet demand, remains optimal. Other parameters will also be given multiple values within $q \in [0.1, 1.5]$, $T \in \{3, 6\}$. Capacity flexibility in the commitment contract is defined as a percentage k so that

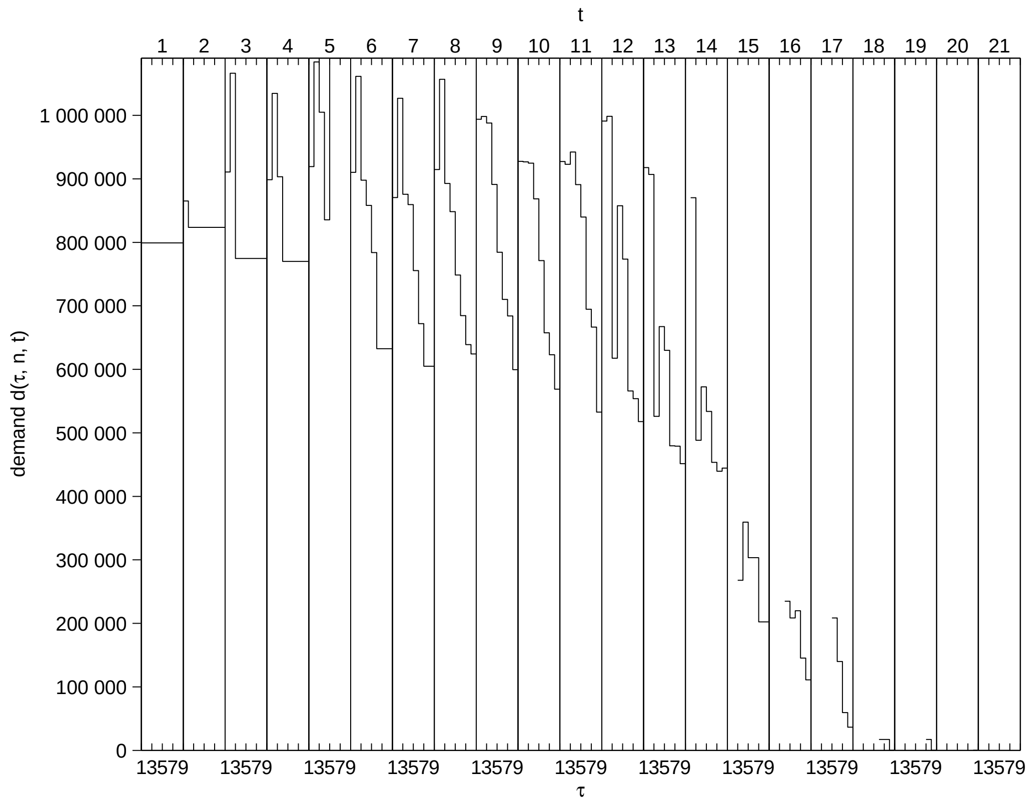

Figure 3.2: An example of demand forecast evolution over monthly iterations.

$K = kQ$. The parameter k is also given different values. Finally, the set of products included in the analysis, P, is given different values to simulate scenarios where the same component is used across multiple products and the demand for all those products can be pooled together.

Chapter 4

Results

The model was created using MATLAB. Creating 1000 demand scenarios took on average 10 seconds on a modern laptop. Calculating costs for 1000 scenarios took on average 0.10 seconds. No special attention was paid on optimization of computational performance in either the scenario creation or cost calculation.

4.1 Demand scenario creation

Past forecasts and realized demands are put to (3.5) to estimate the variance-covariance matrix $\mathbf{\Sigma}$ of the distribution. After deriving the matrix $\mathbf{\Sigma}$, 1000 random forecast update vectors are generated by drawing a random sample from a multivariate normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$. A total of 12 update vectors are applied to the initial plan sequentially until realized demands for $\bar{\tau}$ to $\bar{\tau} + 11$ are available.

An example of the resulting set of demand scenarios is presented in figure 4.1. The frequency distributions of simulated demands are presented on vertical axes which intersect the horizontal axis for the relevant month of the distribution. To improve the readability of the chart, the bars are scaled so that the one representing the highest frequency for a month equals the distance of adjacent vertical axes in length.

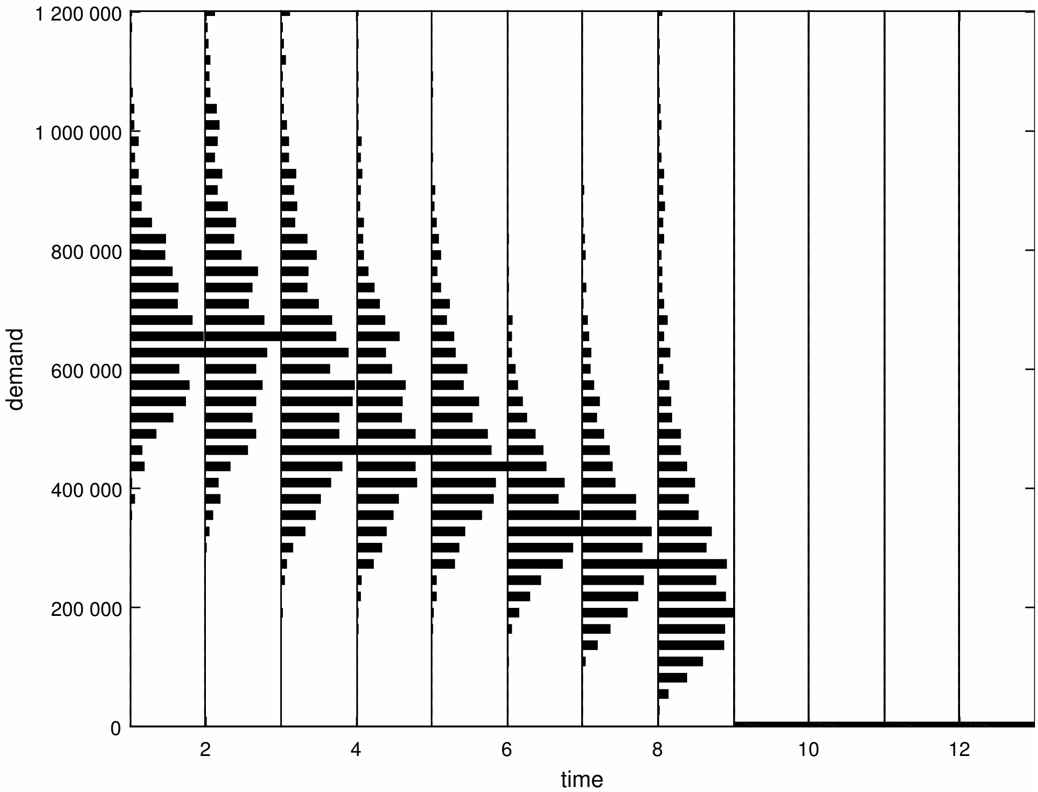

Figure 4.1: Demand distributions for each month using the default parameters.

4.2 Cost savings

Using the default parameters $p(\tau) = 2, d = 0.10, c_h = 0.25, p_c = 0.3, p_u = 3, p_s = 0.4, T = 3, k = 0.2, P = \{6\}$, the model yields results in figure 4.2 where the different metrics are plotted. The results for the default parameters are also included in tables 4.1 and 4.2 on the third row. As can be seen from the results, the maximum occurs between $q = 77\%$ and $q = 81\%$ depending on the metric. The mean savings are at greatest 3.0% with a 5.9% standard deviation. $q = 1$ yields on average a 1.6-percent loss compared to the reference case. Savings can be expected when q is between 69% and 95%. The five metrics follow a similar curve with CVaR 5% as the lowest one, constantly negative as the other VaR and CVaR curves. The deterministic curve reaches d as shown in chapter 3.7.

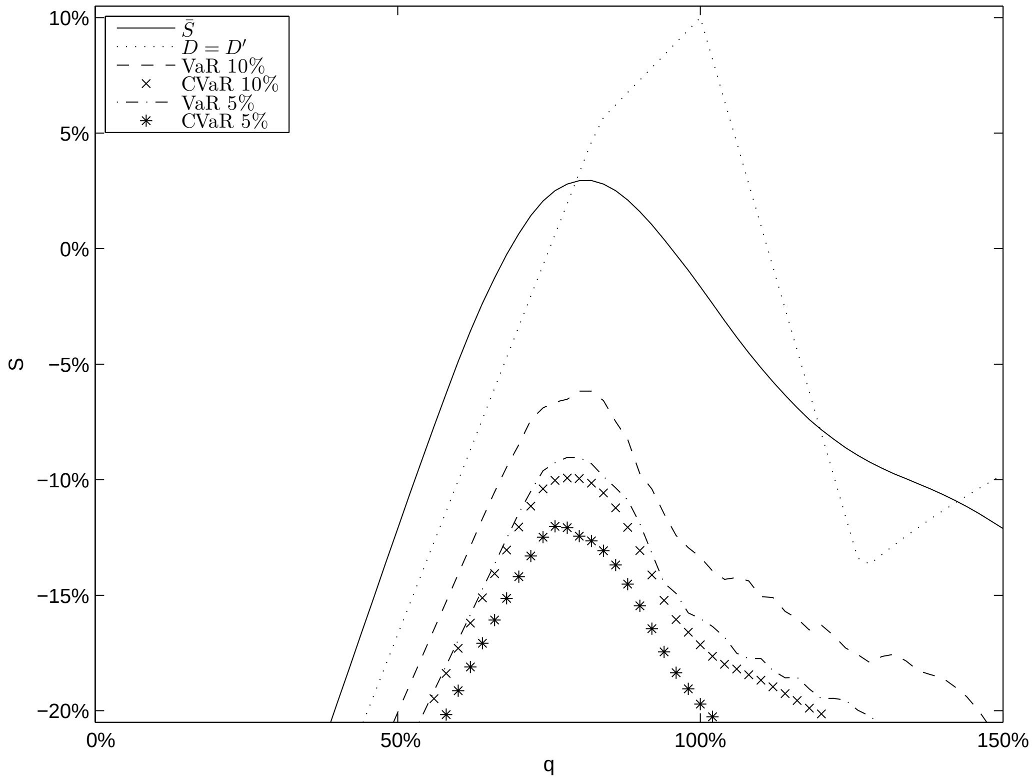

Figure 4.2: Mean savings, deterministic savings, and VaR and CVaR at 5% and 10% using the default parameters.

The distributions for cost savings can be presented in the same way as demand distributions. In figure 4.3, the distributions are plotted at 10% intervals on q-axis. This figure shows that for example when $q = 1$, there is a wide spread of outcomes with similar probabilities. It also helps understand the risk by visualizing the tail for $q = 0.8$ and $q = 1$. This complements the VaR and CVaR metrics presented earlier.

A cost breakdown can also been made for C_A and the different elements of it. In figure 4.4, the cost breakdown for the default problem is presented. The commitment cost increases linearly while the cost of unmet demand decreases as the commitment moves to levels closer to the initial plan. Inventory carrying cost and salvage cost are barely visible with the default

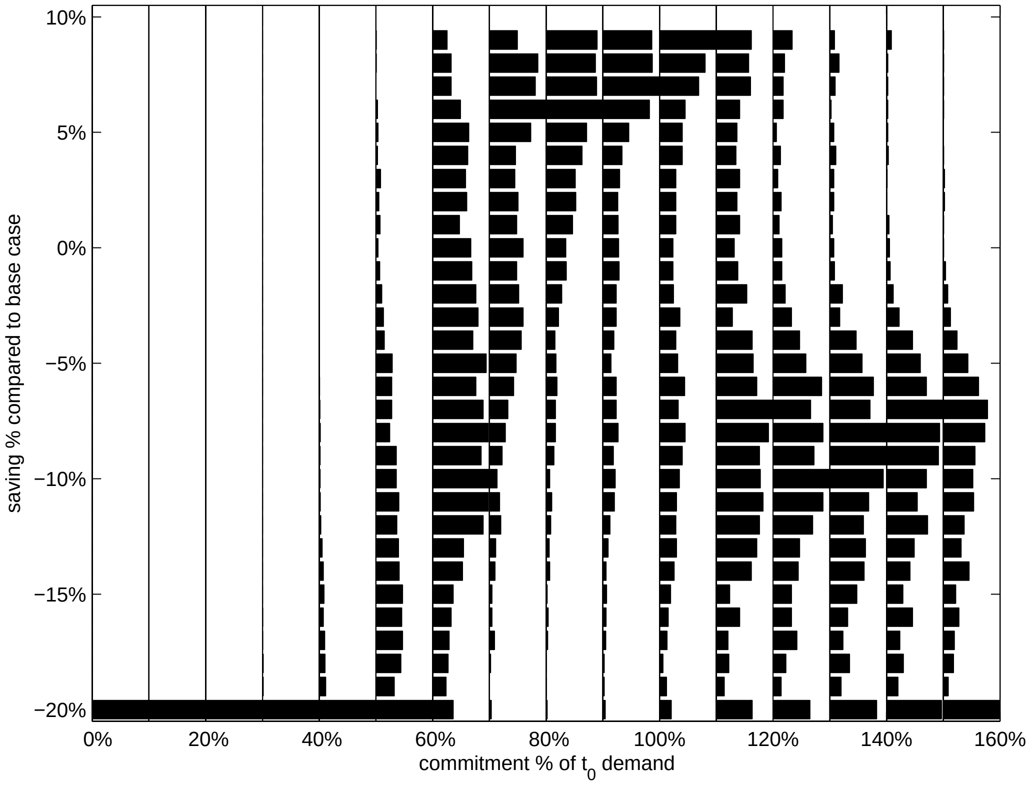

Figure 4.3: Savings distributions using the default parameters.

parameters. However, setting $P = \{16\}$, which represents a product that is closer to the end of its life cycle, yields a similar cost breakdown but with a greater likelyhood of carring excess inventory and ending up with inventory to salvage, as shown in figure 4.5.

Tables 4.1 and 4.2 also contain results for different levels of d. As expected, higher discount makes the commitment contract more attractive for the buyer. The optimal commitment level is higher for greater discount. The same applies to all metrics studied while there is variation between the q^* of different metrics.

There are a number of other changes to the default parameters presented in tables 4.1 and 4.2. Setting a longer commitment period $T = 6$ shifts the optimal commitment quantity to 91%. The mean savings are lower and standard deviation higher indicating that there's a bigger risk of demand

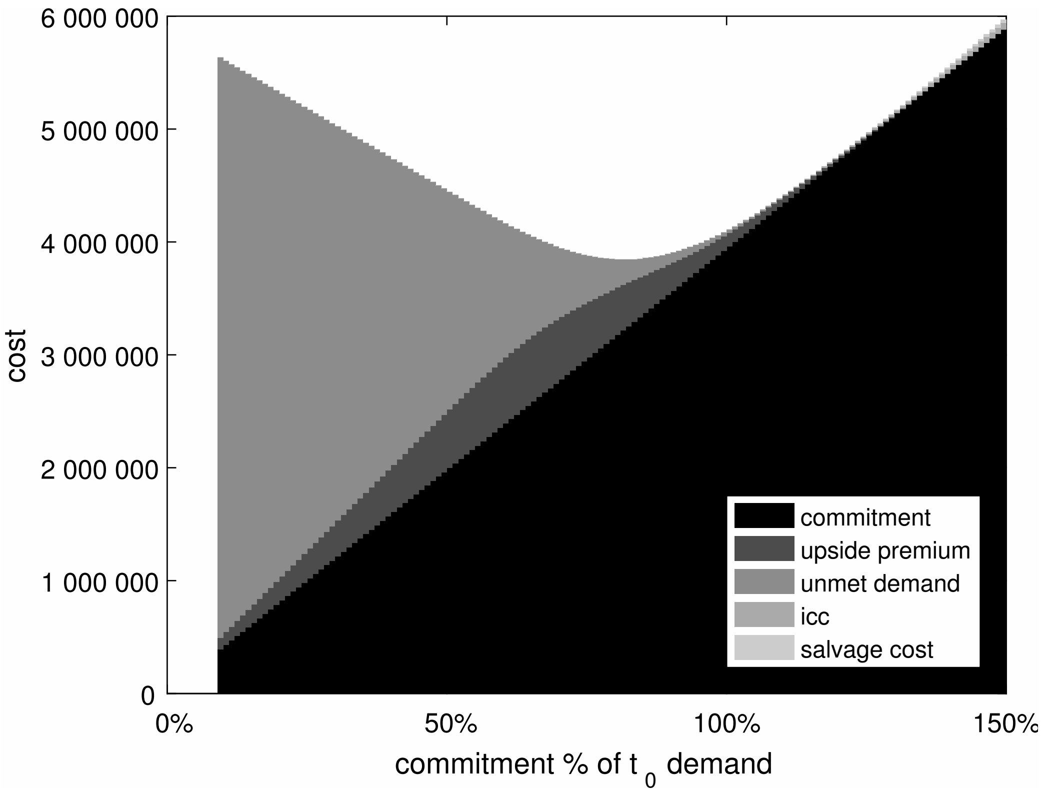

Figure 4.4: Cost breakdown for different cost elements using the default parameters.

changes causing additional cost.

When increasing the level of flexibility available k the metrics quickly saturate to certain values indicating that there are few extreme upside cases in the data. On the other hand, the default penalty for unmet demand refers to a case where the manufacturer may buy the components from another supplier B if maximum capacity is used with supplier A discounting the value of upside flexibility. Quantity commitment with no flexibility, $k = 0$, shows that for default parameters, the commitment is expected to yield only 0.5% savings.

Changing the penalty for unmet demand p_u to values that describe the cost of lost sales changes the game for the manufacturer. Optimal values for q increase as the penalty gets greater. Losing flexibility is expensive at the

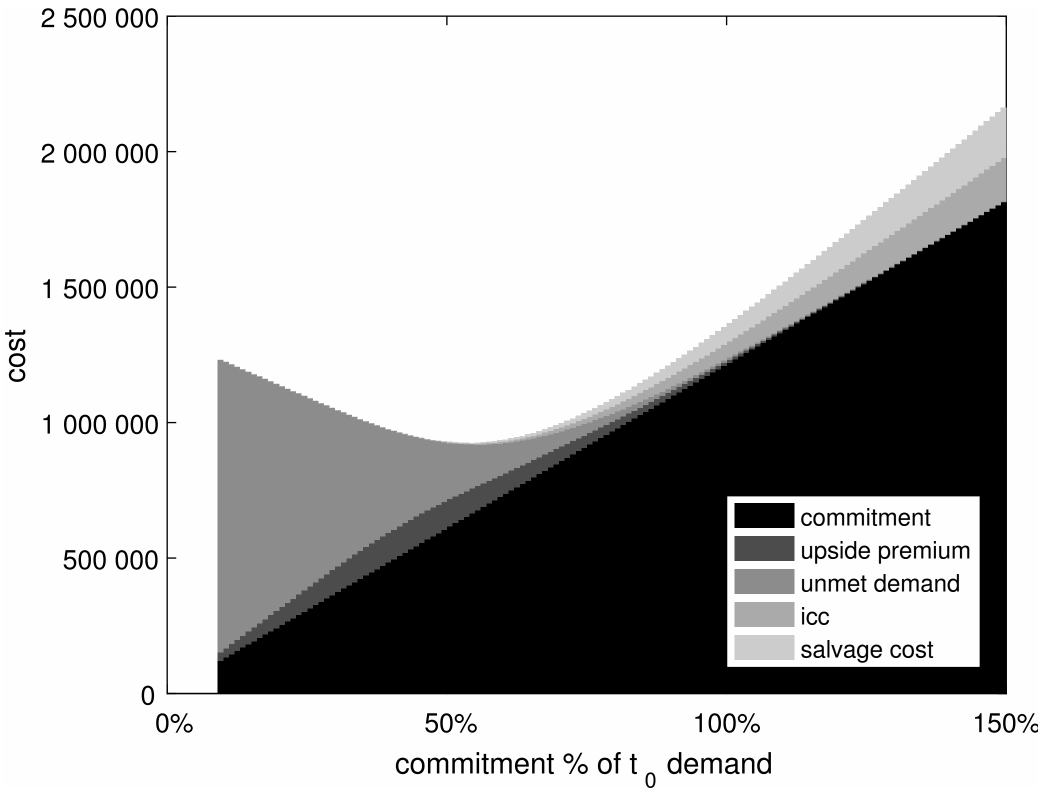

Figure 4.5: Cost breakdown for different cost elements for $P = \{16\}$.

higher end of the scale for p_u and commitment, even though it still comes with an option for a twenty-percent upside, is not expected to yield savings at any commitment level. S for different levels of p_u is plotted in Figure 4.6. As the commitment quantity approaches levels where upsides are non-existent, also the savings curves converge. However, for the greater values of p_u, all of this happens below the break-even line when keeping the rest of the default parameters fixed.

Final variation of the default parameters is in the form of changing the set of products included in the analysis, P. In the default case, $P = \{6\}$, and further scenarios are created by adding products so that $P_2 = \{6, 1\}, P_3 = \{6, 1, 2\}, \ldots, P_{16} = \{1, \ldots, 16\}$. It is notable that while the expected savings are also higher than in the case of $P = \{6\}$, the standard deviation is lower and VaR measures higher when more products are added. This indicates

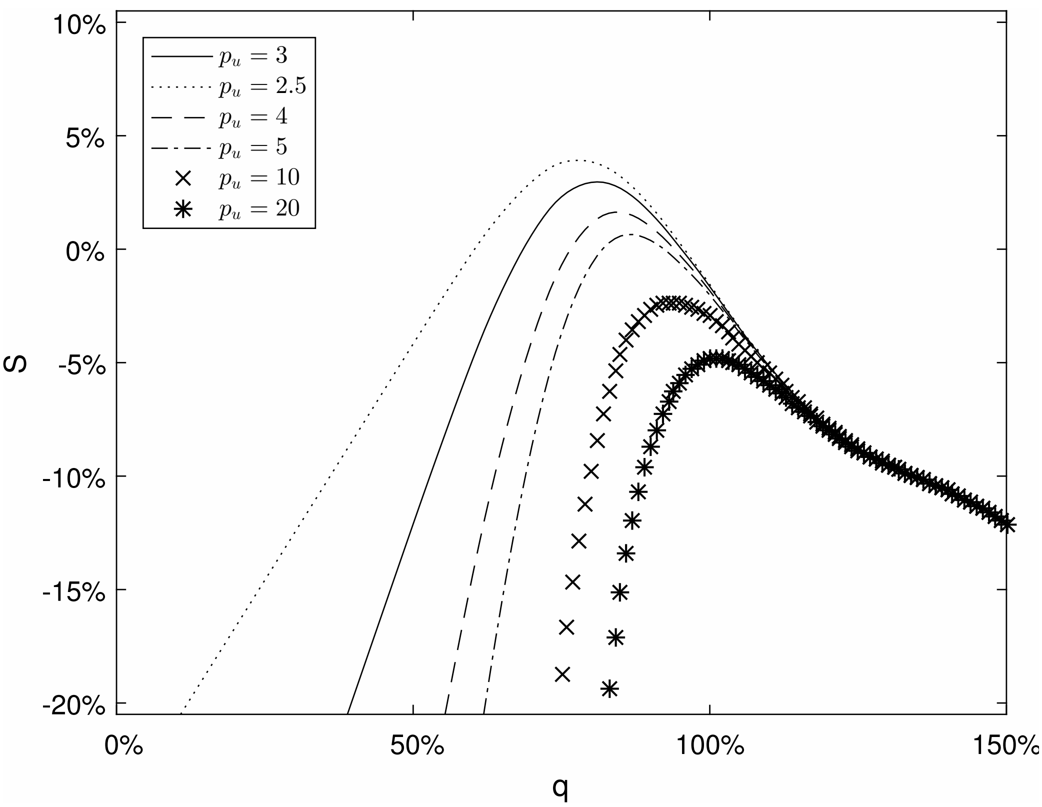

Figure 4.6: Mean savings at different levels of p_u.

that by pooling the demand of all products together, the variations in the demand of a single product do not affect the total savings as much.

Table 4.1: Optima and optimal value; standard deviation for savings; $q = 1$ expected savings; range for expected savings for default parameters and scenarios where one parameter had been changed. The change against default parameters is in the first column.

Change	$q^*, \bar{S}^*$	σ_{S^*}	$\bar{S_{100}}$	q_l, q_s
$d = 0.02$	80%, -5.2%	5.9%	-10.6%	-
$d = 0.05$	80%, -2.1%	5.9%	-7.3%	-
-	81%, 3.0%	5.9%	-1.6%	[69%, 95%]
$d = 0.15$	82%, 8.1%	6.0%	4.0%	[61%, 112%]
$d = 0.2$	83%, 13.2%	5.9%	9.6%	[56%, >150%]
$T = 6$	91%, 2.2%	7.9%	0.6%	[79%, 101%]
$k = 0$	84%, 0.5%	6.3%	-2.4%	[80%, 90%]
$k = 0.4$	78%, 3.9%	5.0%	-1.5%	[64%, 95%]
$k = 0.6$	77%, 4.2%	4.4%	-1.5%	[60%, 95%]
$k = 0.8$	76%, 4.3%	4.2%	-1.5%	[58%, 95%]
$k = 1$	76%, 4.3%	4.1%	-1.5%	[57%, 95%]
$p_u = 2.5$	78%, 3.9%	4.7%	-1.6%	[61%, 95%]
$p_u = 4$	84%, 1.6%	7.8%	-1.8%	[77%, 94%]
$p_u = 5$	87%, 0.6%	9.3%	-2.0%	[82%, 92%]
$p_u = 10$	94%, -2.3%	14.2%	-3.0%	-
$p_u = 20$	101%, -4.8%	18.4%	-4.8%	-
$P = P_2$	64%, 4.5%	5.0%	-10.2%	[53%, 75%]
$P = P_3$	54%, 5.9%	3.8%	-10.6%	[44%, 63%]
$P = P_4$	52%, 5.4%	4.0%	-9.7%	[43%, 62%]
$P = P_5$	54%, 5.0%	4.6%	-9.3%	[45%, 64%]
$P = P_6$	57%, 5.5%	4.0%	-9.6%	[47%, 67%]
$P = P_7$	56%, 5.9%	3.7%	-9.4%	[46%, 66%]
$P = P_8$	59%, 6.6%	3.0%	-10.9%	[48%, 69%]
$P = P_9$	57%, 6.5%	3.2%	-10.9%	[46%, 67%]
$P = P_{10}$	57%, 6.3%	3.3%	-11.2%	[47%, 68%]
$P = P_{11}$	58%, 6.3%	3.4%	-11.4%	[48%, 68%]
$P = P_{12}$	58%, 6.6%	3.2%	-11.5%	[47%, 67%]
$P = P_{13}$	58%, 6.5%	3.1%	-12.0%	[47%, 68%]
$P = P_{14}$	66%, 6.4%	3.3%	-13.4%	[54%, 77%]
$P = P_{15}$	67%, 6.7%	3.0%	-13.7%	[54%, 78%]
$P = P_{16}$	67%, 6.7%	2.9%	-13.3%	[54%, 78%]

Table 4.2: Optima and optimal values for VaR and CVar metrics at 5% and 10% levels. The change against default parameters is in the first column.

Change	$q^*_{\text{VaR 5\%}}$, VaR 5%*	$q^*_{\text{VaR 10\%}}$, VaR 10%*	$q^*_{\text{CVaR 5\%}}$, CVaR 5%*	$q^*_{\text{CVaR 10\%}}$, CVaR 10%*
$d = 0.02$	73%, -16.5%	79%, -13.4%	74%, -19.2%	74%, -17.2%
$d = 0.05$	74%, -13.8%	79%, -10.8%	75%, -16.6%	75%, -14.6%
-	79%, -9.0%	81%, -6.1%	77%, -12.0%	79%, -9.9%
$d = 0.15$	81%, -4.0%	83%, -1.0%	83%, -7.2%	82%, -4.8%
$d = 0.2$	83%, 1.2%	84%, 4.2%	85%, -1.8%	84%, 0.4%
$T = 6$	84%, -11.2%	88%, -7.2%	82%, -15.8%	83%, -12.9%
$k = 0$	83%, -11.2%	85%, -8.4%	83%, -14.1%	83%, -11.9%
$k = 0.4$	73%, -6.6%	78%, -3.2%	74%, -9.8%	75%, -7.5%
$k = 0.6$	71%, -4.0%	76%, -0.3%	72%, -7.6%	72%, -5.0%
$k = 0.8$	70%, -1.8%	76%, -0.2%	70%, -5.3%	70%, -3.2%
$k = 1$	70%, -1.8%	76%, -0.2%	69%, -4.0%	69%, -2.7%
$p_u = 2.5$	71%, -3.9%	77%, -1.9%	70%, -6.2%	71%, -4.7%
$p_u = 4$	88%, -12.9%	87%, -9.8%	90%, -17.4%	88%, -14.5%
$p_u = 5$	93%, -15.2%	87%, -11.7%	99%, -20.4%	93%, -17.5%
$p_u = 10$	101%, -18.1%	98%, -14.5%	114%, -22.6%	113%, -19.8%
$p_u = 20$	115%, -18.8%	106%, -14.9%	123%, -24.4%	120%, -21.1%
$P = P_2$	61%, -5.5%	62%, -2.6%	61%, -8.5%	61%, -6.3%
$P = P_3$	52%, -2.1%	53%, 0.8%	52%, -4.4%	52%, -2.5%
$P = P_4$	51%, -2.4%	50%, -0.3%	51%, -5.4%	51%, -3.4%
$P = P_5$	52%, -4.2%	54%, -1.9%	52%, -7.0%	52%, -5.1%
$P = P_6$	56%, -2.8%	56%, -0.4%	56%, -5.5%	56%, -3.6%
$P = P_7$	55%, -1.8%	55%, 0.5%	55%, -4.2%	55%, -2.4%
$P = P_8$	60%, 0.2%	59%, 2.5%	58%, -2.4%	58%, -0.6%
$P = P_9$	57%, 0.1%	57%, 2.2%	56%, -2.6%	56%, -0.7%
$P = P_{10}$	58%, -0.4%	58%, 1.7%	57%, -3.2%	57%, -1.2%
$P = P_{11}$	57%, -0.8%	58%, 1.5%	57%, -3.5%	57%, -1.6%
$P = P_{12}$	56%, 0.2%	57%, 2.4%	57%, -2.5%	57%, -0.7%
$P = P_{13}$	57%, 0.3%	58%, 2.2%	57%, -2.5%	57%, -0.6%
$P = P_{14}$	65%, 0.5%	65%, 2.0%	64%, -2.6%	65%, -0.8%
$P = P_{15}$	66%, 1.0%	67%, 2.8%	65%, -1.6%	66%, 0.1%
$P = P_{16}$	65%, 1.3%	67%, 3.0%	65%, -1.4%	66%, 0.3%

Chapter 5

Summary

A model for describing uncertainty in customer demand and modeling procurement cost related to two types of procurement contracts was presented in this thesis. The model selection was motivated by a real-life scenario where a manufacturer updates her demand forecasts monthly using a 12-month rolling horizon and purchasing components using a price-only contract. The alternative contract evaluated was a total quantity commitment contract with flexibility. Using a MMFE model for the demand data and a deterministic cost model for procurement cost, the manufacturer's procurement cost was calculated and optimized. This provides the manufacturer with a tool for assessing whether they should enter a quantity commitment contract with the supplier or how they should try to affect the parameters of such contract under negotiation, taking into account that the discount the supplier offers for the commitment may be offset by other cost elements incurred by the contract.

The methods presented in this thesis give the manufacturer visibility on the expected savings depending on the parameters as well as shows how the demand uncertainty affects the savings. It is good to remember that the risk level increases with the commitment which is why also measures describing the risk (VaR and CVaR) were studied. One approach that the manufacturer could adopt is minimizing risk exposure rather than maximizing expected savings. In some cases, this can mean giving up expected savings, but that may be justified for instance for a small company that cannot sustain a major financial loss. It is also valuable for any decision maker to understand the distribution of the outcome beyond the aggregated statistical indicators. Visualizing the distributions and making them available to the decision makers by using modern business intelligence solutions can help contingency planning by explicitly reminding that demand uncertainty is not always a minor inaccuracy of the forecast but extreme scenarios also exist.

What makes this approach more attractive for industry practitioners is the modularity of the solution. Demand uncertainty is experienced by many companies but their cost models may differ considerably. Hence, it is valuable to change either model without having to re-create the whole model. The use of the demand scenario asset is not only limited to the procurement context. In parallel, the production planning department could validate that they have enough capacity or necessary mitigation plans to support the demand scenarios. Further, the model could for example consider demand uncertainty in geographical dimension to support logistics network planning. In all of these cases, Monte Carlo simulation is a useful method for optimizing because it is easy to use and makes it easy to combine any models that are used for demand uncertainty and the optimization.

Further research be conducted by studying the same situation from the supplier's point-of-view to understand whether the contracts that are attractive to the manufacturer are also attractive to the supplier, considering that she also has suppliers facing the same uncertainty. The study could be extended to cover multiple tiers of upstream supply chain to understand how different procurement contracts affect the cost of each tier. A commitment contract downstream may actually help coordinate the supply chain to avoid the bullwhip effect (see e.g. Lee, Padmanabhan, and Whang 1997). In this thesis, the lead times were not modelled into the system because the focus was purely on the manufacturer's point-of-view, but the impact of those, among with some of the other simplifications made, may be an interesting topic for further research and worthy of study in a multi-tier supply chain model as well.

References

Anupindi, R. and Y. Bassok (1999). Supply contracts with quantity commitments and stochastic demand. In S. Tayur, R. Ganeshan, and M. Magazine (Eds.), *Quantitative Models for Supply Chain Management*, pp. 197–232. Boston, MA, United States: Springer.

Ballou, R. (2004). *Business Logistics Management: Supply Chain Management. Planning, Organizing, and Controlling the Supply Chain* (5 ed.). Harlow, United Kingdom: Pearson Education Limited.

Bassok, Y. and R. Anupindi (1997). Analysis of supply contracts with total minimum commitment. *IIE Transactions 29*(5), 373–381.

Bassok, Y. and R. Anupindi (2008). Analysis of supply contracts with commitments and flexibility. *Naval Research Logistics (NRL) 55*(5), 459–477.

Burnetas, A., S. M. Gilbert, and C. E. Smith (2007). Quantity discounts in single-period supply contracts with asymmetric demand information. *IIE Transactions 39*(5), 465–479.

Burrows, P. (2003). Cisco's comeback. *Business Week*, 42–48. November 24.

Cachon, G. P. (2003). Supply chain coordination with contracts. *Handbooks in Operations Research and Management Science 11*, 229–340.

Cachon, G. P. and M. A. Lariviere (2001). Contracting to assure supply: How to share demand forecasts in a supply chain. *Management Science 47*(5), 629–646.

Cachon, G. P. and P. H. Zipkin (1999). Competitive and cooperative inventory policies in a two-stage supply chain. *Management Science 45*(7), 936–953.

Callioni, G., X. de Montgros, R. Slagmulder, L. N. Van Wassenhove, and L. Wright (2005). Inventory-driven costs. *Harvard Business Review 83*(3), 135–141.

Cattani, K. and W. Hausman (2000). Why are forecast updates often disappointing? *Manufacturing & Service Operations Management 2*(2), 119–127.

Chand, S., V. N. Hsu, and S. Sethi (2002). Forecast, solution, and rolling horizons in operations management problems: A classified bibliography. *Manufacturing & Service Operations Management 4*(1), 25–43.

Chen, F. Y. and D. Krass (2001). Analysis of supply contracts with minimum total order quantity commitments and non-stationary demands. *European Journal of Operational Research 131*(2), 309–323.

Ellram, L. M. (1995). Total cost of ownership: An analysis approach for purchasing. *International Journal of Physical Distribution & Logistics Management 25*(8), 4–23.

Erhun, F. and S. Tayur (2003). Enterprise-wide optimization of total landed cost at a grocery retailer. *Operations Research 51*(3), 343–353.

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

Gaur, V., S. Kesavan, A. Raman, and M. L. Fisher (2007). Estimating demand uncertainty using judgmental forecasts. *Manufacturing & Service Operations Management 9*(4), 480–491.

Graves, S. C., H. C. Meal, S. Dasu, and Y. Qui (1986). Two-stage production planning in a dynamic environment. In S. Axsäter, C. A. Schneeweiss, and E. Silver (Eds.), *Multi-stage Production Planning and*

Inventory Control, Lecture notes in economics and mathematical systems, pp. 9–43. Berlin, Germany: Springer.

Gupta, A. and C. D. Maranas (2003). Managing demand uncertainty in supply chain planning. *Computers & Chemical Engineering 27*(8), 1219–1227.

Hausman, W. H. (1969). Sequential decision problems: A model to exploit existing forecasters. *Management Science 16*(2), B–93.

Hausman, W. H. and R. Peterson (1972). Multiproduct production scheduling for style goods with limited capacity, forecast revisions and terminal delivery. *Management Science 18*(7), 370–383.

Heath, D. C. and P. L. Jackson (1994). Modeling the evolution of demand forecasts with application to safety stock analysis in production/distribution systems. *IIE Transactions 26*(3), 17–30.

Kaminsky, P. and J. M. Swaminathan (2004). Effective heuristics for capacitated production planning with multiperiod production and demand with forecast band refinement. *Manufacturing & Service Operations Management 6*(2), 184–194.

Lariviere, M. A. (1999). Supply chain contracting and coordination with stochastic demand. In S. Tayur, R. Ganeshan, and M. Magazine (Eds.), *Quantitative Models for Supply Chain Management*, pp. 233–268. Boston, MA, United States: Springer.

Lee, H. L., V. Padmanabhan, and S. Whang (1997). Information distortion in a supply chain: The bullwhip effect. *Management science 43*(4), 546–558.

Martínez-de Albéniz, V. and D. Simchi-Levi (2005). A portfolio approach to procurement contracts. *Production and Operations Management 14*(1), 90–114.

Nagali, V., J. Hwang, D. Sanghera, M. Gaskins, M. Pridgen, T. Thurston, P. Mackenroth, D. Branvold, P. Scholler, and G. Shoemaker (2008).

Procurement risk management (prm) at Hewlett-Packard company. *Interfaces 38*(1), 51–60.

Porteus, E. L. (1990). Stochastic inventory theory. *Handbooks in Operations Research and Management Science 2*, 605–652.

Ren, Z. J., M. A. Cohen, T. H. Ho, and C. Terwiesch (2010). Information sharing in a long-term supply chain relationship: The role of customer review strategy. *Operations Research 58*(1), 81–93.

Research In Motion (2011). Research In Motion Announces Third Quarter Provision Related To PlayBook Inventory and Confirms Commitment to Tablet Market; Provides Update to Q3 and Fiscal 2012 Guidance. `http://press.blackberry.com/financial/2011/pressrelease-5267.html` (accessed February 3rd, 2014).

Shapiro, J. (2001). *Modeling the Supply Chain.* Belmont, CA, United States: Brooks/Cole-Thomson Learning.

Sodhi, M. and C. S. Tang (2010). Determining supply requirement in the sales-and-operations-planning (S&OP) process under demand uncertainty: A stochastic programming formulation and a spreadsheet implementation. *Journal of the Operational Research Society 62*(3), 526–536.

Sodhi, M. S. (2005). Managing demand risk in tactical supply chain planning for a global consumer electronics company. *Production and Operations Management 14*(1), 69–79.

Stevenson, W. (2009). *Operations Management* (10 ed.). Operations and decision sciences. New York, NY, United States: McGraw-Hill Irwin.

Szegö, G. (2002). Measures of risk. *Journal of Banking & Finance 26*(7), 1253–1272.

Tang, C. S. (2006). Perspectives in supply chain risk management. *International Journal of Production Economics 103*(2), 451–488.

Terwiesch, C., Z. J. Ren, T. H. Ho, and M. A. Cohen (2005). An empirical analysis of forecast sharing in the semiconductor equipment supply chain. *Management Science 51*(2), 208–220.

Tsay, A. A., S. Nahmias, and N. Agrawal (1999). Modeling supply chain contracts: A review. In S. Tayur, R. Ganeshan, and M. Magazine (Eds.), *Quantitative Models for Supply Chain Management*, pp. 299–336. Boston, MA, United States: Springer.

Van Landeghem, H. and H. Vanmaele (2002). Robust planning: A new paradigm for demand chain planning. *Journal of Operations Management 20*(6), 769–783.

Van Weele, A. (2009). *Purchasing and Supply Chain Management: Analysis, Strategy, Planning and Practice* (5 ed.). London, United Kingdom: Cengage Learning.

Waters, C. (2003). *Logistics: An Introduction to Supply Chain Management*. Basingstoke, United Kingdom: Palgrave MacMillan.

Wisner, J. (2011). *Principles of Supply Chain Management: A Balanced Approach* (3 ed.). Mason, OH, United States: South-Western.

Wu, D. and D. L. Olson (2008). Supply chain risk, simulation, and vendor selection. *International Journal of Production Economics 114*(2), 646–655.

Wu, S. D., B. Aytac, R. T. Berger, and C. A. Armbruster (2006). Managing short life cycle technology products for Agere Systems. *Interfaces 36*(3), 234–247.