



Norwegian University of
Science and Technology

Allocating Sales in the Farming of Atlantic Salmon

Maximizing Profits Under Uncertainty

Martin Bergan Hæreid

Industrial Economics and Technology Management

Submission date: July 2011

Supervisor: Asgeir Tomasgard, IØT

Co-supervisor: Kristin Tolstad Uggen, IØT
Peter Schütz, IØT

Norwegian University of Science and Technology

Department of Industrial Economics and Technology Management

MASTERKONTRAKT

- uttak av masteroppgave

1. Studentens personalia

Etternavn, fornavn Hæreid, Martin Bergan	Fødselsdato 25. apr 1986
E-post martin@hareid.com	Telefon 99032082

2. Studieopplysninger

Fakultet Fakultet for Samfunnsvitenskap og teknologiledelse	
Institutt Institutt for industriell økonomi og teknologiledelse	
Studieprogram Industriell økonomi og teknologiledelse	Hovedprofil Anvendt økonomi og optimering
E-post martin@hareid.com	Telefon 99032082

3. Masteroppgave

Oppstartsdato 17. jan 2011	Innleveringsfrist 13. jun 2011
Oppgavens (foreløpige) tittel Allocating Sales in the Farming of Atlantic Salmon Maximizing Profits Under Uncertainty	
Oppgavetekst/Problembeskrivelse The thesis presents an application of multistage stochastic programming to the contract portfolio planning problem facing a Norwegian salmon farmer. A model aims to maximize profit by allocating sales between different types of contracts and sales in the spot market. Decisions are made considering estimates of future salmon prices and the uncertainty associated with the growth of salmon, along with the many constraints that together define salmon aquaculture. The objective is to develop a decision support model, and the main topics to be included are as follows: <ol style="list-style-type: none"> 1. How to deal with non-linear relationships in the formulation 2. An assessment of uncertainty relevant to the model 3. An implementation of the model in Xpress 4. A test case along with a discussion of the results 	
Hovedveileder ved institutt Professor Tomasgard Asgeir	Biveileder(e) ved institutt Kristin Tolstad Uggen, Peter Schütz
Merknader 1 uke ekstra p.g.a påske.	

4. Underskrift

Student: Jeg erklærer herved at jeg har satt meg inn i gjeldende bestemmelser for mastergradsstudiet og at jeg oppfyller kravene for adgang til å påbegynne oppgaven, herunder eventuelle praksiskrav.

Partene er gjort kjent med avtalens vilkår, samt kapitlene i studiehåndboken om generelle regler og aktuell studieplan for masterstudiet.

Trondheim, 17.01.2011

Sted og dato

Martin B. Høeneid

Student



Hovedveileder

Originalen oppbevares på fakultetet. Kopi av avtalen sendes til instituttet og studenten.

Preface

This master thesis is the final step in achieving a Master of Science degree at the Norwegian University of Science and Technology (NTNU). The degree specialization is Applied Economics and Optimization at the Department of Industrial Economics and Technology Management. The thesis has been written in cooperation with SINTEF Technology and Society and SINTEF Fisheries and Aquaculture, and is motivated by their ongoing work with Marine Harvest ASA.

I have received much appreciated help and guidance in completing this thesis, both from co-students and teachers at NTNU, as well as several industry contacts. I would like to thank my supervisor, professor at NTNU and senior scientist at SINTEF Technology and Society, Asgeir Tomasgard, for both patience and constructive discussions. Peter Schütz and research fellow Marte Fodstad at SINTEF Technology and Society have provided invaluable help in formulating and implementing the stochastic model. Additionally, I would like to thank Eivind Osnes and colleagues at Marine Harvest and Lars Liabø and colleagues at Kontali Analyse. Due to limited literature on salmon aquaculture, their cooperation has been a key part of gaining an in-depth insight into the industry and obtaining realistic input data for the model. My work has greatly benefited from all of the mentioned above.

Trondheim, June 1, 2011

Martin Bergan Hæreid

Abstract

Salmon farmers face an uncertain production environment and considerable price volatility, making planning a vital success criteria. This thesis describes the sources of uncertainty that are most important when planning sales, and demonstrates how this uncertainty can be taken into consideration by the use of stochastic programming.

The basis for this thesis is the tactical planning problem of deciding when to harvest salmon and how to allocate sales between available contracts and sales in the spot market. Uncertainty relevant to the planning problem is described, and a multistage stochastic model that maximizes profit is proposed. The goal of the model is to provide salmon producers with a tool that can aid them in making profitable decisions regarding harvesting and future sales, by taking into account the uncertainty associated with biomass and price development.

The model is implemented in three versions; a deterministic model, a two-stage stochastic model, and a multistage stochastic model. The implemented models are somewhat simplified, the most important simplification being that price is assumed deterministic in the stochastic models. This is done in order to make the stochastic models computationally tractable for a personal computer. All three models are written in Mosel, implemented in Xpress-IVE, and solved by the Xpress Optimizer.

The implemented models are applied to Marine Harvest Region Mid, illustrating how the models can be used to solve a realistic salmon sales planning problem. In addition, a quantitative assessment of the gains from implementing a stochastic solution is demonstrated. The results obtained show that using the two-stage stochastic model provides almost no additional value over the deterministic model. For the multistage stochastic model, this value is substantially higher, though still marginal, largely due to the simplifications made in the implementation. Based on the simplifications made in the implemented models, possible extensions to the thesis are suggested.

Contents

List of Figures	vii
List of Tables	viii
1 Introduction	1
2 Atlantic Salmon Aquaculture	3
2.1 Industry Background	3
2.2 Production in Seawater	4
2.3 The Norwegian Regulatory Framework	5
2.4 Uncertainty in Salmon Farming	6
2.4.1 Biomass Development	6
2.4.2 The Market for Atlantic Salmon	11
2.4.3 Other Sources of Uncertainty	15
3 Stochastic Programming	17
3.1 Dealing with Uncertainty	17
3.2 Recourse Problems	18
3.2.1 Two-stage Recourse	18
3.2.2 Recourse with Multiple Stages	21
3.2.3 Scenario Trees	22
3.2.4 Recourse Models and Rolling Horizon	23
3.3 Evaluation of Recourse Models	24
3.3.1 The Value of the Stochastic Solution & the Expected Value of Perfect Information	25
3.3.2 VSS in Multistage Models	27
4 Model Introduction	31
4.1 Model Objective	31
4.2 Model Scope	32
4.3 Using the Model	33
4.4 Model Input	34
5 A Multistage Stochastic Model	37

5.1	Modelling Growth and Decision Variables	37
5.2	Formal Definitions	41
5.3	Model Formulation	43
5.3.1	Objective Function	43
5.3.2	Constraints	43
5.3.3	The Complete Model	48
6	Model Application	51
6.1	The Models	51
6.1.1	Simplifications	52
6.1.2	Implementation of the Models	52
6.2	Data Set	53
6.2.1	Marine Harvest Region Mid	53
6.2.2	Resolution and Planning Horizon	54
6.2.3	Biomass Development	54
6.2.4	Scenarios	54
6.2.5	Prices and Contracts	57
6.2.6	Additional Data	58
6.3	Two-stage Stochastic Model	59
6.4	Multistage Stochastic Model	60
7	Results and Discussion	65
7.1	Main Results	65
7.1.1	EV, RP, WS and EEV	65
7.1.2	Summarized Results	67
7.2	Detailed Results	68
7.3	Evaluation and Analysis	74
7.3.1	VSS and EVPI	74
7.3.2	Spot Purchase Price and Contract Price Markups	75
8	Concluding Remarks	79

List of Figures

2.1	Yearly mean runoff to the Sognefjord, 1985 to 2001	8
2.2	Deviation from the in-sample monthly mean temperature, location in Sør-Trøndelag, 1998 to 2006	8
2.3	Monthly mean temperature and the in-sample relative standard deviation, location in Sør-Trøndelag, 1998 to 2006	9
2.4	Monthly growth for a 5 kg salmon, location in Sør-Trøndelag, 1998 to 2006	10
2.5	Loss in production of salmon, 2000 to 2009	11
2.6	Escaped salmon, 2000 to 2009	12
2.7	Average operating margin, 1986 to 2009	13
2.8	Weekly price of Atlantic salmon, 1995 to 2010	15
3.1	The decision process for a two-stage recourse model	19
3.2	A simple scenario tree	23
3.3	A set of scenario problems and non-anticipativity constraints	24
3.4	Relation between VSS and EVPI	26
6.1	Xpress Product Suite	52
6.2	Growth scenarios	57
6.3	Problem instance 1: 1997	59
6.4	Problem instance 2: 2001	59
6.5	Problem instance 3: 2010	60
6.6	The scenario tree for the two-stage problem	61
6.7	The scenario tree for the multistage problem	62
6.8	Scenarios in the multistage problem	63
7.1	Detailed results for problem instance 1	69
7.2	Detailed results for problem instance 2	70
7.3	Detailed results for problem instance 3	71

List of Tables

6.1	Locations in Marine Harvest Region Mid	55
6.2	Growth model	56
7.1	Results for the two-stage problem	67
7.2	Results for the multistage problem	67
7.3	First stage decisions in problem instance 1	72
7.4	First stage decisions in problem instance 2	73
7.5	First stage decisions in problem instance 3	73
7.6	Evaluation of the TS model	74
7.7	Evaluation of the MS model	74
7.8	Results of varying the spot purchase price markup	75
7.9	Spot purchase price markup evaluation	75
7.10	Results of varying the contract price markup	76
7.11	Contract price markup evaluation	76

Chapter 1

Introduction

Farming of Atlantic salmon has grown to become one of Norway's most important industries during the last decades. The expectations are high as to what the industry can achieve in the years to come, especially in light of the decreasing production outputs in the Norwegian petroleum sector. Even as the salmon aquaculture industry has matured, it is still characterized by large variations in supply and varying margins, largely due to the biological nature of salmon farming and the resulting risks. Dealing with these risks requires understanding the uncertainty present, and is an essential part of successful planning in salmon farming. This thesis is motivated by the potential gains of utilizing stochastic optimization in planning, and the absence of work done within this area for the salmon farming industry.

The scope of this thesis is limited to dealing with the seawater part of the salmon farming value chain. This part begins when young salmon are released in the sea, and lasts until the salmon are fully grown and ready to be sold. While in the ocean, the salmon are subject to varying growth conditions, potential disease outbreaks, and escape, all of which introduce uncertainty in the production process. Then, when the salmon are ready to be slaughtered and sold, the salmon farmer faces a salmon price that exhibits significant volatility. The first goal of this thesis is to study and describe the uncertainty that is most relevant when planning harvesting and sales, and present stochastic programming as a tool for dealing with this uncertainty. Second, a multistage stochastic model is proposed that considers both uncertainty in salmon growth and price development. The aim of the model is to provide salmon producers with a tool that can aid them in making profitable decisions regarding harvesting and future sales. Finally, the model is implemented and applied to Marine Harvest Region Mid, illustrating how the model can be used to solve a realistic salmon sales planning problem.

The thesis consists of eight chapters, and is summarized in the following. After this introduction, chapter 2 explains the basics of salmon aquaculture, providing an overview of relevant aspects of the industry along with a discussion of the

uncertainty relevant to the planning of harvesting and sales. Chapter 3 gives an introduction to stochastic programming and how stochastic models can be evaluated, covering theory that is relevant for dealing with the uncertainty present in the salmon industry. A multistage stochastic model is introduced in chapter 4, along with an explanation of how the model is to be used. A detailed description of the model is given in chapter 5. Chapter 6 presents an application of the model to Marine Harvest Region Mid, while chapter 7 gives the results of the implementation. Chapter 8 concludes with a discussion of natural extensions to the model and the possibilities for future work.

Throughout the thesis, salmon implies Atlantic salmon, and aquaculture is short for Atlantic salmon aquaculture. Unless otherwise noted, all weights are in whole fish equivalents, WFE, which are measured in metric tons. The terms *salmon producer* and *salmon farmer* are used unambiguously, and both refer to a company operating in the seawater part of the value chain.

Chapter 2

Atlantic Salmon Aquaculture

When studying the economics of salmon aquaculture and topics related to optimization it is natural to begin with an introduction of today's salmon industry. This chapter introduces the basics of salmon aquaculture, providing relevant background information for the model presented in chapter 5. Section 2.1 gives a brief overview of the industry, section 2.2 explains the production process in seawater, section 2.3 treats relevant aspects of the Norwegian regulatory framework, while the uncertainty involved with salmon production is discussed in section 2.4.

2.1 Industry Background

Even though the art of salmon farming has roots over 200 years back, it was not until modern aquaculture techniques were developed in the late 1960s that an industry began to develop. Since then, advances in technology have led to the supply of farmed salmon surpassing that of wild salmon, with Norway leading the way. Today, about two thirds of the world's total supply of salmon species is from aquaculture (Kontali Analyse, 2010). But even with an increase in worldwide production of Atlantic salmon of more than 600% since 1990 (Marine Harvest, 2010c), salmon species still only make up roughly 2.3% of global seafood supply. In 2009, a total of 1 467 000 tons of Atlantic salmon were harvested worldwide, with approximately 855 700 tons originating in Norway (Trollvik, 2010). This was the highest quantity ever produced in Norway, and preliminary estimates of 925 000 tons for 2010 indicate continued growth driven by increasing demand (Trollvik, 2010).

Since salmon aquaculture requires specific environmental conditions, the indus-

try is dominated by North America, UK, Chile and Norway. These regions all have sheltered coastlines, which combined with suitable seawater temperatures and appropriate governmental legislation facilitates for successful salmon farming. Together, the four regions represent approximately 95 % of total worldwide harvest. In addition to being the world's by far largest producer, Norway's high production to national demand ratio also makes the country the world's largest exporter. In 2010, Norway's seafood export totaled 53.8 billion NOK, whereof 31.4 billion NOK was Atlantic salmon. Salmon produced in Norway is primarily supplied to the EU. This supply, along with supply from Chile to USA, Canada to USA, and UK to EU, make up what historically has been the four main trade flows. This situation has during the last few years adjusted slightly due to growing demand from Asia, fluctuating supply from Chile, and an increase in secondary processing, which together have resulted in a more globalized market. But, with no indication of processed or frozen products dominating fresh salmon in the near future, the majority of supply will most likely continue to follow the traditional trade flows (Marine Harvest, 2010c).

2.2 Production in Seawater

After 6 to 18 months in freshwater, salmon undergo a process called smoltification. Smoltification enables the fish to survive in seawater, facilitating the transition from land based production in fresh water to marine farms along the coast. Salmon are at this stage referred to as smolt, and usually weigh from 60 to 150 grams at the time of release. Once in seawater, salmon are farmed in net cages containing anywhere from approximately 30 000 to 200 000 fish, depending on size and layout of the production facility (Cermaq ASA, 2011). Net cages are rectangular or round, and are made of either plastic or steel. Rectangular cages most commonly range in size from 10 by 10 to 40 by 40 meters, while circular cages can be anywhere from 40 to 200 meters in circumference (AKVA Group, 2010).

Feeding of the salmon is normally done using automatic feeding systems that follow advanced veterinary plans ensuring optimal conditions for fish health and growth. As the fish grow, sorting may be required, either because only fish of a certain size are to be harvested, or because lowering size variations in a cage increases feed intake and prevents harmful hierarchies from developing among the fish (The Research Council of Norway, 2005). Sorting can be performed using well boats or on-site sorting equipment.

After 12 to 18 months in the sea, the salmon are ready to be slaughtered. Fish can either be slaughtered immediately on-site, or transported with a well boat to the nearest slaughtering facility. Here, the fish are gutted, cleansed, sorted, filleted and packed. The salmon are either sold as whole fish or processed further and sold as value added products (Cermaq ASA, 2011). Once a site is harvested, the location must be fallowed for 2 to 6 months before it is ready to be used by the

next generation of salmon (FHL Havbruk, 2005). This gives a total production time for Atlantic Salmon of approximately 24 to 36 months (Marine Harvest, 2010c).

2.3 The Norwegian Regulatory Framework

Environmental challenges involved with farming salmon in seawater has necessitated stringent governmental regulation. This was most recently illustrated by the lack of regulation in Chile, and the vast problems and decline in supply that followed (Barros, 2011). In Norway, legislation is exercised by the Ministry of Fisheries and Coastal Affairs, the Directorate of Fisheries, and The Norwegian Food Safety Authority. This section will focus on the regulations that are most relevant to the model presented in chapter 5, and will therefore not cover the extensive restrictions concerning feeding and slaughtering routines, control of fish welfare, required production routines, or other restrictions that apply before or after farming in seawater. The reader is referred to *Akvakulturdriftsforskriften* (2008) for a complete list of the regulations concerning operation of aquaculture facilities.

Licensing has been practiced in Norway since 1973, and is the main tool for ensuring sustainable development in the industry. Today's regulatory framework was established in 2005, and is described in *Laksetildelingsforskriften* (2004). In 2010, a total of approximately 105 companies were operating 960 licenses in Norway (Sandberg, 2010). Licenses are awarded by the Ministry of Fisheries, while governing is left to the Directorate of Fisheries. In most of Norway, the biomass limit per license is normally 780 tons live weight, while in the northernmost parts, the limit is 900 tons. A license is valid within a specified region, and the sum of licenses in a region gives the region's maximum allowable biomass, or MAB. Within a region, one license can be utilized in up to four sites, while two licenses together can be utilized in up to six sites. In addition, each site has a site specific MAB, which can be exploited using one or more licenses.

The most recent development in the Norwegian regulatory framework is *Luseforskriften* of 2009. This regulation gives The Norwegian Food Safety Authority the authority to establish regions in which fallowing must be conducted simultaneously for all facilities (Ministry of Fisheries and Coastal Affairs, 2009). The purpose is to avoid contamination between different locations and generations of fish. Currently, the regions are only established in certain parts of the country.

In addition to regulating the production in seawater, Norwegian authorities have also set a limit for the maximum market share a company may have. To prevent companies from having disproportionate amounts of market power, no more than 25 % of the available licenses can be in the hands of the same owner. Also, if purchasing a license results in the buyer controlling more than 15 % of the total allowable biomass in Norway, the company must seek permission from The

Ministry of Fisheries and Coastal Affairs (Marine Harvest, 2010a). Currently, the market share restriction is only relevant for Marine Harvest ASA, which after rapid expansion and several acquisitions over the last years controls approximately 25% of the market.

2.4 Uncertainty in Salmon Farming

Producers in general face two main types of risk: production risk and price risk. Production risk refers to the effects of uncertainty in the output of production, while price risk describes the effects of uncertainty in the revenue obtained from the quantity produced (Just and Pope, 1978; Sandmo, 1971). In salmon aquaculture, both of these risks are present, production risk through the uncertainty in biomass development, and price risk, through the uncertainty associated with entering contracts and trading fish in the spot market. In a long term perspective, additional uncertainty is also present, mainly surrounding the strategic conditions under which salmon producers operate. Note that this section focuses primarily on explaining and measuring the sources of uncertainty, without a detailed discussion of the resulting risk. For a further discussion regarding the risks in salmon farming, the reader is referred to Kumbhakar and Tveterås (2003), Bergfjord (2009), and Tveterås (1999).

2.4.1 Biomass Development

There are two main causes to the uncertainty in biomass development. Firstly, there is the uncertainty associated with salmon growth rates. This includes both the uncertainty in overall growth due to uncertain growth conditions, and the uncertainty in the growth distribution within a biomass caused by varying growth rates among fish of the same weight. Secondly, loss of fish, either due to mortality or escape, can drastically reduce biomass. Due to the presence of this uncertainty combined with the salmon's high sensitivity to its marine environment, it is argued that salmon aquaculture is more risky than traditional land-based livestock production. In addition, farmers can more easily control important biophysical variables in land-based production (Tveterås, 1999).

Growth Rates

Important biological parameters affecting salmon health are oxygen concentration, salinity, pH, ammonia and carbon dioxide content, lighting conditions, feed, sea temperature, and more. Of these factors, lighting and feed can to a large degree be controlled by the producer, while the first four factors mentioned are usually within acceptable values (Institute of Marine Research, 2009). This leaves temperature as the main source of the uncertainty in determining growth rates. Marine farms are typically located in sheltered surroundings in fjords or in coastal areas, where the variations in water temperature can be significant. Since salmon

are cold blooded, and therefore extremely sensitive to changes in temperature, these variations in temperature are important to understand.

Water temperatures along the Norwegian coast are determined by a combination of oceanic and local influences. In oceanic waters, temperatures follow a relatively certain seasonal pattern, largely due to the high thermal capacity of water. The high thermal capacity implies that it takes much longer to heat or cool water compared to for example air, which has approximately one fourth the thermal capacity. This gives oceanic water a relatively good “memory”, meaning that the temperature in a period is highly dependent on the temperature in the previous period, assuming periods of limited length. But in near shore waters where marine farms are located, the effect of oceanic water is limited. Here, the water temperature is mainly determined by local weather conditions and water exchange between coastal waters and fjords (Sætre, 2011).

The water exchange between coastal waters and fjords is mainly influenced by freshwater runoff and by differences in sea level and water density. The mean annual freshwater runoff from Norway is round 12 000 cubic meters per second, or 400 cubic kilometers per year, and all this flows through the fjords. Freshwater induced exchange mainly occurs in the upper layers of the water, and is therefore of great relevance to salmon growth. The runoff follows a clear seasonal pattern, but with large inter-annual fluctuations, as shown in figure 2.1. The figure shows the annual runoff to the longest Norwegian fjord, the Sognefjord, in the period 1985 to 2001. In regards to differences in sea level between coastal waters and fjords, these are mainly caused by winds along the coast and tidal variations. The tidal variations caused by the semidiurnal tide have the greatest effect, due to their rapid fluctuations and relatively large tidal differences. The tidal variations increase as one moves north along coast, and can be as much as 2.7 meters in the northernmost parts of Norway. Finally, there are the effects of differences in density between coastal waters and fjords. Density fluctuations in coastal waters can generate horizontal coast/fjord pressure differences, which induce flows in or out of the fjord. The density distribution of coastal waters may be changed by either advection of new water masses with different properties, or by wind-induced coastal upwelling or downwelling (Sætre, 2007).

The most important aspects of local weather conditions are whether it is sunny or cloudy and the amount of wind. These factors often have an immediate effect on the temperature in the uppermost layer of water where the salmon are kept. Since future weather conditions are hard to predict except for the immediate future, uncertainty is introduced.

In sum, the effects described above result in coastal waters, as opposed to ocean waters, having almost no intertemporal temperature dependency when disregarding the seasonal pattern. This is illustrated by figure 2.2. The figure shows the deviation from the in-sample monthly mean temperature for a salmon farm in Sør-Trøndelag for the period 1998 to 2006, with no apparent correlation from one period to the next. The result is that foreseeing future deviations from the expected seasonal pattern is very difficult in areas where the marine farms are




Figure 2.1: Yearly mean runoff to the Sognefjord, 1985 to 2001 (Sætre, 2007)

located.




Figure 2.2: Deviation from the in-sample monthly mean temperature for a salmon farm in Sør-Trøndelag, 1998 to 2006 (Marine Harvest, 2011)

Figure 2.3 displays the monthly mean temperature for the same salmon farm for the years 1998 to 2006, showing considerable variations in temperature from the seasonal pattern. Each year is represented by a colored line, labeled by the left vertical axis. Also shown is the in-sample relative standard deviation, shown as the dashed line, and labeled by the right vertical axis. Notice that the uncertainty is especially large from June to September. The effect of these temperature variations on salmon growth are illustrated in figure 2.4, which shows the monthly growth for a 5 kg salmon for the same time period. In addition to the uncertainty discussed above, it is also worth mentioning the variations in growth rates among salmon that are the same size. These variations can be approximated by a normal distribution with a variation coefficient of 22% (Marine Harvest, 2011).




Figure 2.3: Monthly mean temperature and the in-sample relative standard deviation for a salmon farm in Sør-Trøndelag, 1998 to 2006 (Marine Harvest, 2011)

Loss in Production

In addition to the escape of salmon, the main reasons for losses occurring while the salmon are in the ocean are mortality of young fish and disease outbreaks. Approximately 80% of fish mortality occurs before the salmon are 0.5 kg, due to either deformities, injuries from transportation or release, or fish not coping with the transition to seawater (Fauske, 2011). In addition to causing mortality, disease affects biomass development by causing reduced appetite and situations where premature slaughtering is necessary. Therefore, much effort is put into minimizing the concentration of bacteria, viruses, and toxic algae. Some of the most important diseases include Infectious Pancreatic Necrosis (IPN), Pancreas Disease (PD), Heart and Skeletal Muscle Inflammation (HSMI), Infectious Salmon Anaemia (ISA), Salmonid Rickettsial Septicaemia (SRS), Gill Disease




Figure 2.4: Monthly growth for a 5 kg salmon in a salmon farm in Sør-Trøndelag, 1998 to 2006 (Skretting, 2011; Marine Harvest, 2011)

(GD), and sea lice (Marine Harvest, 2010c). The number of salmon lost in production (excluding escape) reported to the Directorate of Fisheries over the last 10 years are shown in figure 2.5, compared to salmon sales in the same period. Note that the figure only shows the number of deaths, not the negative effects that disease may have in terms of reduced growth and premature slaughtering.

The main reasons for farmed salmon escaping are winter storms, propeller damage, and wear and tear on equipment. Fish escaping through holes in nets account for approximately 65% of the reported cases (AKVA Group, 2010). In recent years, better management of these problems has led to a reduced number of escaped salmon, which contrasts with the increased number of salmon produced. The number of escaped salmon reported to the Directorate of Fisheries over the last 10 years are shown in figure 2.6, compared to total salmon sales in the same period. Large escapes (> 10 000 fish) account for only 19% of all incidents, but amount to 91% of the number of escaped fish (Østen, 2010). This means that even though the number for escaped salmon is small in comparison with losses due to mortality, the consequence of an escape can be far greater for the unlucky producer.

Consequences of the Uncertainty in Biomass Development

If not handled carefully, the uncertainties in biomass development discussed above can have costly consequences. Higher growth than expected may necessitate unplanned harvesting in order to satisfy MAB restrictions, resulting in biomass which must be sold in the spot market. This increases the exposure to the




Figure 2.5: Loss in production of salmon, 2000 to 2009 (Directorate of Fisheries, 2010)

price risk discussed in following sections. Another possible consequence of the salmon growing faster than expected is that the salmon reach the delivery weight earlier than planned. This might require renegotiation of contract terms or the purchase of a replacement product in the spot market. In cases where the biomass development is lower than expected, a producer may be forced to purchase fish in the spot market in order to fulfill contract commitments. This can be costly since the producer might not be able to buy a replacement product in time due to limited availability, or since a premium above the spot price must be paid when buying the product from a competing producer. If the product is not delivered, non-refundable costs (transportation, booked slaughtering capacity, etc.) and costs associated with loss of customer goodwill may apply (Asche, 2010). In addition, since poor growth conditions and disease outbreaks often affect the entire industry, chances are high that the shortage applies to the majority of producers. The result may be a decrease in market supply combined with increase in producer demand (to fulfill contracts), thus putting pressure on prices and making purchasing fish in the spot market undesirable.

2.4.2 The Market for Atlantic Salmon

Understanding the dynamics of the market for salmon is essential for any company involved in the salmon industry. According to surveys done by Bergfjord (2009), Norwegian fish farming companies consider future salmon prices to be the most important source of risk. This is not surprising, considering the turbu-




Figure 2.6: Escaped salmon, 2000 to 2009 (Directorate of Fisheries, 2010)

lent economic environment that has characterized the industry during its rapid development over last decades. A substantial variation in margins has led to frequent bankruptcies, often associated with periods of persistently low salmon prices (Bergfjord, 2009). Figure 2.7 shows the yearly average operating margin (operating result/operating income) for the period 1986 to 2009. Even as the industry has matured, future salmon prices remain hard to foresee. This is largely because of the unpredictable supply resulting from the biological nature of salmon farming, and the almost perfect negative correlation between the supply of salmon and the salmon market price (Marine Harvest, 2010c).

Supply

In line with economic theory, supply and demand determine the price in a competitive market. In the salmon industry, the volatility in this price can in part be explained by looking at the price elasticity of supply. As most of the farmed salmon is marketed fresh and is thereby perishable, salmon produced in one period has to be consumed in the same period. In the short term, the production level is difficult and expensive to adjust, as the planning/production cycle is approximately three years long. This, combined with government regulations limiting capacity, results in a short run price supply elasticity that is close to zero (Andersen and Tveterås, 2008)¹. At the same time, short term supply is greatly affected by uncertain biological factors and the resulting development in biomass, as discussed above. The result is a short term supply that is both in-

¹The study showed that a 1% increase in the sales price only induced about a 0.05% increase in supply.




Figure 2.7: Average operating margin, 1986 to 2009 (Directorate of Fisheries, 2010)

elastic and uncertain, hereby introducing the possibilities for both periods of over and undersupply, causing prices to fluctuate. In addition, even though a variety of sizes and quality creates different product segments, salmon can be considered a relatively homogeneous product. This, along with the development of international trade, has created a worldwide market where the price in a region is not only sensitive to changes in local supply, but also governed by changes in the supply in other countries. This was recently illustrated by the outbreak of the ISA virus in Chile in 2007, where the reduction in supply created an imbalance between global supply and demand.

In the long run, the situation changes substantially as supply becomes elastic. Asche and Salvanes (1997) estimated a long-run supply elasticity of about 1.5 using annual aggregated data, and Asche and Tveterås (2007) reported the same, deriving the results from a cost function. In Andersen and Tveterås (2008), the difference in price responsiveness of salmon supply in the short and long run is argued to play a large part in explaining the observed cyclical profitability in the salmon farming industry. The reader is referred to Andersen and Tveterås (2008) for a more detailed discussion on the reasons and consequences of the elasticity of supply being larger in the long run than in the short run.

Demand

With future salmon prices also being a function of future demand, it is natural that the results in Bergfjord (2009) show that future demand is perceived as one

of the most importance sources of risk by salmon producers. A number of studies investigating the demand for salmon have shown that the demand elasticity is on a decreasing trend, with the current demand elasticity in Europe being approximately -1.0 (Asche and Bjørndal, 2005). This of course varies between countries, but a common pattern is that fresh salmon seems to be more elastic than frozen, and that salmon can be substituted by other product types, but not by other fish species (Asche and Salvanes, 1998).

The main forces driving salmon demand are factors such as seasonality, changes in consumer preferences and welfare, along with market trends and product awareness. The increasing demand for salmon in recent years is largely credited to the general increase in interest for fish and other seafood products due to the health and nutrition benefits that these products provide. For salmon, the positive effects of omega-3 are especially important. At the same time, the industry faces challenges surrounding the effects of salmon farming on the environment, and questions have been raised regarding whether today's industry represents a sustainable management of natural resources. Here the main issues are sources of feed, fish welfare, and the effects of salmon farming on wild salmon and the marine habitats surrounding farms.

Price

Resulting from the discussion above is a salmon price that exhibits significant volatility. Figure 2.8 shows the weekly price of Atlantic salmon for the period 1995 to 2010. The prices shown in the figure are for fresh, head-on, gutted, superior salmon, packaged and delivered from slaughtering, also referred to as HOG (head-on gutted). The average price for Norwegian whole salmon the last decade has been approximately 25 NOK/kg HOG, with peaks at 45 NOK/kg and lows at 15 NOK/kg. Studies done by Alnæs and Skagen (2009) show that the annualized volatility at times during this period approached 50%. For a discussion of the characteristics of salmon price volatility, see Oglend and Sikveland (2008). In addition, Asche and Guttormsen (2002) provide insight into how the biological production cycle can cause patterns in the relative prices between the different sizes throughout the year.

For salmon farmers wishing to secure future volumes and prices, the traditional method has been to enter long term contracts with processing companies, wholesalers or supermarkets. Now, with the recent establishment of Fish Pool, financial tools present new possibilities for reducing the exposure to price risk. Fish Pool is the first and only organized marketplace for trading derivatives with salmon as the underlying asset. Trading financial salmon contracts provides salmon farmers with an alternative to entering physical contracts, with the additional benefits that a purely financial market entails. With the longest contracts stretching nearly three years, it is theoretically possible for salmon farmers to secure all or part of their earnings at a contractual price. But, Fish Pool is still in an early stage of development, and investigations done by Alnæs and Skagen (2009) show that limited liquidity prevents the futures, forwards, and options available




Figure 2.8: Weekly price of Atlantic salmon, 1995 to 2010 (Norwegian Seafood Federation, 2010)

from being fully efficient, thereby limiting the possibilities for creating a perfect hedge.

2.4.3 Other Sources of Uncertainty

In addition to the risk associated with uncertain future prices and biomass development, the exploratory study done by Bergfjord (2009) shows that Norwegian fish farming companies consider institutional risk factors to be of great importance. In the study, market regulation, area access, changes to the licensing system, environmental regulations, taxation, certification systems, and animal health regulations are mentioned as the most important institutional risk factors. Governmental restrictions, like quotas and tolls set by the EU on Norwegian salmon, are most always enforced after a prior warning, and will thereby give producers the chance to adapt. But in the long term, these changes can in any case be vital to overall yield in the industry. This is clearly illustrated by the modernization of the industry in Norway, which has undergone dramatic structural and legislative changes since companies in 1991 for the first time were allowed to own more than one license. The introduction of maximum allowable biomass in 2005, along with the more recent lice zone restrictions, have been important steps in bringing the industry closer towards a viable state of operation. It is predicted that the strict governmental involvement will continue, and that the coming years will without doubt bring additional regulations that again will change the conditions under which salmon farmers operate.

Other sources of uncertainty worth mentioning are the price of feed, future exchange rates, health concerns, and changing natural conditions. As the salmon industry continues to globalize, currency exchange rates become ever more important, especially for export dependent countries like Norway. In regards to changing natural conditions, there is an uncertainty associated with the effects of a general increase in sea temperature in Norway due to global warming. Salmon growth rates are highest when the temperature is between 13 and 17 degrees Celsius, which for Norway implies that a limited increase in sea temperature will have a positive effect on productivity. But if the temperature increases to the point where it approaches the limits of what is suitable for the salmon, mortality rates will increase, and the salmon industry might have to move farther north to ensure suitable conditions (Lorentzen, 2010).

Chapter 3

Stochastic Programming

This chapter provides an introduction to stochastic optimization, serving as a framework for later discussions on how to deal with the uncertainty present in the salmon industry. Section 3.1 introduces the role of uncertainty in optimization problems, followed by a presentation of recourse models in section 3.2. Section 3.3 concludes with a discussion of how stochastic models can be evaluated. The goal of the chapter is not to provide a complete coverage of the topics above, but rather present the reader with an overview of important terms and concepts, which will ease the understanding of the model presented in the next chapter. For a more thorough introduction to stochastic programming in general, the reader is referred to Kall and Wallace (1994) or Birge and Louveaux (1997). It is assumed that the reader is familiar with basic principals in optimization and mathematical programming.

3.1 Dealing with Uncertainty

Uncertainty is present in most real life decision problems, and can result from prices, demand, costs, weather, technology and more. Taking this uncertainty into account is an essential part of successfully modelling a decision problem, and is a topic that has received considerable attention in operations research literature. In dealing with uncertainty, there are several possible approaches, the simplest of which is a *deterministic problem formulation*. A deterministic approach does not include the uncertainty directly in the model, but instead relies on either careful determination of the input parameters or thorough analysis of the solution, using methods such as sensitivity analysis, what-if analysis, or scenario analysis (Midthun, 2010). In sensitivity analysis, the robustness of a solution is analyzed by varying the uncertain parameters, and noting the resulting changes in the optimal solution. Alternatively, a scenario analysis can be performed, where a number of scenarios are created based on possible realizations of

uncertain parameters. For each scenario, the deterministic problem is solved, and by combining and analyzing the results, an optimal solution can be found.

Alternatively, uncertainty can be taken into account by being included directly in the problem formulation. This is referred to as *stochastic programming*, which in some cases can be a far superior approach to dealing with uncertainty compared to the deterministic methods described above (Wallace, 2003). The goal of stochastic programming is to find some policy that is feasible for all (or almost all) of the possible data instances, while maximizing or minimizing the expectation of some function of random variables (Philpott, 2011). Put in another way, stochastic programming can be viewed as a tool for finding all the explicit and implicit options worth paying for in a decision problem (Midthun, 2010). A fundamental assumption is that the probability distribution of the random variables, representing the possible outcomes and their respective probabilities, is either known or can be estimated.

The two topics that have received the majority of attention in stochastic programming literature are *recourse models* and *chance constrained programming*. The concept of recourse is the basis for the stochastic model presented in the next chapter, and the remainder of this chapter is therefore devoted to this topic. For an introduction to chance constrained programming, the reader is referred to Birge and Louveaux (1997) or Prékopa (1995).

3.2 Recourse Problems

In recourse models, decisions are classified according to when they have to be made in relation to resolution of uncertainty in the decision problem. This adds value to the model by utilizing the flexibility of postponing decisions until information about uncertainties is revealed (Sen and Higle, 1999). By introducing *stages* based on when new information is available, an opportunity is given to adapt a solution to a specific outcome observed, hence the term recourse. Recourse models are therefore always presented as models in which there are two or more stages, allowing for the exploitation of relevant information that becomes available during the planing horizon (Higle, 2005).

3.2.1 Two-stage Recourse

The most widely applied and studied stochastic programming models are two-stage linear programs with recourse (Philpott, 2011). These models are appropriate when some decisions must be fixed before information relevant to the uncertainties is available, while the remaining decisions can be delayed until after. More specifically, the decision maker takes some action in the first stage, after which a random event occurs affecting the outcome of the first-stage decision. A recourse decision can then be made in the second stage, compensating

for any negative effects or exploiting positive effects that resulted from the first-stage decision. The optimal solution presented is a single first-stage decision, along with a collection of recourse decisions defining which actions to be taken in the second-stage in response to each possible outcome of the random variables (Philpott, 2011).

Figure 3.1 illustrates this decision process, which is divided in two by ω , the realization of uncertainties. x denotes the *first stage* or *here and now* decision, which is taken before information regarding uncertainties in the problem is revealed. y , on the other hand, is determined after observations regarding ω have been made, and is known as the *second stage* or *recourse variable* (Higle, 2005). The goal of the two-stage recourse model is to identify a first-stage decision x that leaves y well positioned against all possible realizations of the uncertainty (Midthun, 2010).




Figure 3.1: The decision process for a two-stage recourse model (van der Vlerk, 2011)

The structure of a recourse problem can have important implications for feasibility and computational demand. Depending on the a problem's properties, it can be classified as having either general, simple, fixed or complete recourse. Only general recourse will be discussed here, and the reader is referred to Birge and Louveaux (1997) for a presentation of the remaining topics.

Implicit Representation of a Recourse Problem

Using the notation presented in Higle (2005), a *general recourse model* can be stated as follows:

$$\begin{aligned}
\min \quad & cx + E[h(x, \tilde{\omega})] & (3.2.1) \\
\text{s.t.} \quad & Ax \geq B \\
& x \geq 0
\end{aligned}$$

where

$$\begin{aligned}
h(x, \omega) = \min \quad & q_\omega y & (3.2.2) \\
\text{s.t.} \quad & W_\omega y \geq r_\omega - T_\omega x \\
& y \geq 0
\end{aligned}$$

This formulation is generally referred to as an *implicit representation* of the stochastic problem (Birge and Louveaux, 1997), also known as *compact form* or a *node formulation*. Here, problem 3.2.1 is the *first stage problem*, and problem 3.2.2 is known as the *second-stage problem*, the *subproblem*, or the *recourse problem*. $\tilde{\omega}$ is a random variable defined on a probability space $(\Omega, \mathcal{A}, \mathcal{P})$, and the outcome of $\tilde{\omega}$ is represented by a set of scenarios, $\omega \in \Omega$. A *scenario* is defined by Hige (2005) as one specific, complete, realization of the stochastic elements that might appear during the course of the problem. The subscript ω denotes that an element might vary with scenario, and illustrates how y_ω is scenario dependent, while x is not. The term $E[h(x, \tilde{\omega})]$ in the first stage problem is referred to as the *value function* or *recourse function* (Birge and Louveaux, 1997). One of the advantages of presenting a recourse problem in the implicit form is that the information process in the problem is clearly represented, as well as the resulting properties of the first and second stage decisions (Hige, 2005).

Extensive Representation of a Recourse Problem

An alternative to representing a recourse problem in the implicit form is the *extensive/full form* (Birge and Louveaux, 1997), also known as a *scenario formulation*:

$$\min \sum_{\omega \in \Omega} (cx_\omega + g_\omega y_\omega) p_\omega \quad (3.2.3)$$

$$\begin{aligned}
\text{s.t.} \quad & T_\omega x_\omega + W_\omega y_\omega \geq r_\omega \\
& x_\omega - x = 0, \quad \omega \in \Omega & (3.2.4) \\
& x_\omega, y_\omega \geq 0
\end{aligned}$$

Here, the concept of recourse is not as apparent as for the implicit representation. Instead of dividing decisions between the first and second stage, and only including the recourse variables explicitly in a second-stage problem, the problem

is formulated as a set of subproblems, one for each possible scenario. Each subproblem, also called a *scenario problem*, is associated with a particular scenario and may be looked upon as a deterministic optimization problem. The decision variables are modelled as if they are permitted to depend on the specific scenario encountered, with additional constraints ensuring that the information structure associated with the decision process is honored (Higle, 2005). If the number of scenarios used to represent the uncertainty is finite, the resulting problem is referred to as *deterministic equivalent problem* (Philpott, 2011). Hence, with the set of scenarios Ω being finite, equations 3.2.3 and 3.2.4 give an example of a deterministic equivalent problem formulated in extensive form.

Two main changes from the implicit to the extensive form are worth explaining in more detail. Firstly, the objective in equation 3.2.3 is reformulated. But, since for each scenario $\omega \in \Omega$, $p_\omega = \mathcal{P}\{\tilde{\omega} = \omega\}$, the formulation is equivalent to the expected value calculated in the value function in equation 3.2.1. Secondly, the formulation contains the additional constraints given in equation 3.2.4. These constraints are referred to as *non-anticipativity constraints*, and are included to enforce that decisions taken in different scenarios are consistent with the information available in each stage. In a two-stage model, this might seem trivial, as y_ω should obviously be allowed to vary from scenario to scenario, while x_ω should be equal for all scenarios. If this was not the case, and the value of x_ω was allowed to vary freely in response to each specific scenario, the information structure in the problem would be violated, since the value of x_ω should reflect that the available information in the first stage is the same for all scenarios. The importance of non-anticipativity constraints will become more apparent when they are discussed in a multistage setting in the following sections.

Note that the formulation used in equation 3.2.4 is just one of numerous ways in which non-anticipativity constraints can be implemented. Also note that it is not actually necessary to separately denote first-stage and recourse variables in the implicit representation, since this is enforced by the non-anticipativity constraints. First-stage and recourse variables are denoted separately here only to illustrate to which variables the non-anticipativity constraints apply.

3.2.2 Recourse with Multiple Stages

The extensive representation of a recourse problem presented in the previous section can easily be generalized to allow for additional stages. This can be useful, since in many real life situations, new information is revealed at several points in time in the future, necessitating a *multistage recourse model*. Continuing to adopt the notation used in Higle (2005), the multistage problem can be represented as follows:

$$\min \sum_{\omega \in \Omega} p_{\omega} \sum_{t \in \mathcal{T}} c_{\omega}^t x_{\omega}^t \quad (3.2.5)$$

$$\text{s.t.} \quad \sum_{j=1}^t A_{\omega}^{tj} x_{\omega}^j = b_{\omega}^t, \quad t \in T, \omega \in \Omega \quad (3.2.6)$$

$$\{x_{\omega}^t\}_{\omega \in \Omega, t \in T} \in \mathcal{N} \quad (3.2.7)$$

Here, first-stage and recourse variables are no longer denoted separately, in accordance with the discussion at the end of section 3.2.1. c_{ω}^t represents the objective function coefficients corresponding to scenario ω , and the scenario constraints are represented as multistage constraints in equation 3.2.6. In the same manner as for the two stage implicit representation, the relationship between decisions and the information process in the problem is ensured by non-anticipativity constraints, given in equation 3.2.7. Instead of formulating the constraints explicitly, \mathcal{N} is used to denote the set of non-anticipative solutions.

3.2.3 Scenario Trees

The use of scenario trees can be helpful in understanding the characteristics of a recourse problem and the role of non-anticipativity constraints in extensive formulations. Hight (2005) defines a scenario tree as a structured distributional representation of the stochastic elements in a problem, and the manner in which these elements may evolve during the problem's planning horizon. In general, nodes in the scenario tree represent possible states of the problem where decisions can be made, and uncertainty is resolved where there are at least two branches leading out from a single node. Scenarios, also referred to as scenario problems, are defined as specific paths from the root node to a leaf node. Periods are used to denote time, while stages, typically consisting of several time periods, are defined in regards to when new information is available. In practice, decisions are only made when new information becomes known, or in other words, in nodes that are the result of a branching (Midthun, 2007). A node after a branching is therefore the first node in a stage that includes all nodes until the next branching.

A simple scenario tree is shown in figure 3.2, along with some of the notation introduced above. The root node corresponds to the initial stage, where no specific information regarding the random variables is available. Dashed lines separate the three stages in the tree, while the set of leaf nodes indicate that there are 6 possible scenarios in the problem. There are a total of 6 periods; stages 1, 2 and 3 containing periods [1], [2,3], and [4,5,6], respectively. In each stage, all uncertainty in the stage is resolved in the first period, and the remaining time periods in the stage can therefore be considered deterministic. This implies that all decisions in stages 1, 2 and 3 can be taken in periods 1, 2 and 4, respectively.




Figure 3.2: A simple scenario tree

A second tree is shown in figure 3.3. This tree is meant to illustrate the difference between representing a recourse problem in implicit and extensive form, as well as the importance of non-anticipativity constraints in extensive formulations. The clear difference from figure 3.2 is that in extensive form, the scenario tree consists of a set of individual scenario problems, one for each possible scenario, where each subproblem can be dealt with as if it was deterministic. This results in an increased number of variables (nodes), clearly demonstrating the increased computational demand introduced by using an extensive formulation. Also apparent from figure 3.3 is that the information process in the tree no longer is enforced through parent-child node relationships. Therefore, non-anticipativity constraints are included as ellipsoids, representing that all decisions taken in nodes inside an ellipsoid must be equal. This ensures that in nodes that share the same history of information, the same decisions are made, resulting in the solutions obtained being implementable (Sen and Hige, 1999).

Finally, it is worth noting that a scenario tree is a discrete representation of the problem. This implies that for a recourse problem with continuously distributed random parameters, a discretization of the underlying probability distribution of the uncertain parameters is required (Birge and Louveaux, 1997).

3.2.4 Recourse Models and Rolling Horizon

Rolling horizon decision making is a common business practice for making decisions in problems with multiple stages, and illustrates how recourse models can be utilized. The term *horizon* refers to the number of periods in the future for which the forecast is made. In short, the method consists of repeatedly making

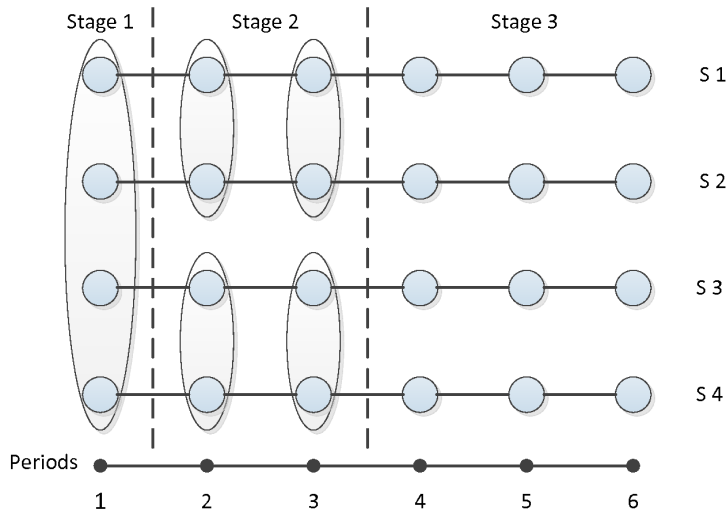


Figure 3.3: A set of scenario problems and non-anticipativity constraints

the most immediate (first-stage) decisions based on forecasts of relevant information concerning future periods. After these decisions are made and as time advances and new information becomes available, a new set of decisions (the second stage decisions) become the most immediate. Before these decisions are made, forecasts for additional periods in the future may be required, and existing forecasts may be revised or updated to reflect new information. This procedure repeats, justifying the term rolling horizon decision making (Sethi and Sorger, 1991).

Using a rolling horizon is often applied to problems where planning continues indefinitely. In Baker (1977), a study of the effectiveness of rolling horizon decision making suggests that there are two principal reasons why finite horizon models might be appropriate for decision making in problems with an infinite horizon. Firstly, decisions must for practical reasons be based on limited information about the future, since forecasting may be both costly and difficult. Secondly, the forecasts for the remote future tend to be unreliable and therefore of limited value. Overall, the study concludes that assuming a finite horizon in combination with rolling horizon decision making is quite efficient.

3.3 Evaluation of Recourse Models

Stochastic models have a reputation for being computationally demanding, and often requiring specific solution methods. It can therefore be useful to have tools for evaluating whether using a stochastic model is necessary, or if it is sufficient to use for example a deterministic approach where the effort instead is aimed

at determining uncertain parameters. The tools for evaluating stochastic models can be especially valuable when the model is going to be used repeatedly and the computational burden is high (Wallace, 2003). Two methods of evaluation are presented in this section: the value of the stochastic solution and the expected value of perfect information. Both of these measures will be used to evaluate the models presented in chapter 6, and since the goal of these models is to maximize profit, the notation below is made in regards to a maximization problem.

3.3.1 The Value of the Stochastic Solution & the Expected Value of Perfect Information

The expected value of perfect information (EVPI) is defined by Birge and Louveaux (1997) as a measure of the maximum amount a decision maker would be ready to pay in return for complete (and accurate) information about the future, thereby removing all uncertainty. The value of the stochastic solution (VSS), on the other hand, measures the expected value of using a stochastic model over a deterministic model. In order to explain how the two measures are calculated, some notation must be introduced.

When all random variables in the problem are replaced by their expected values, one obtains the expected value problem or mean value problem, where the solution is known in relevant literature as the *expected value* (EV) solution. The expected value of using the EV solution when uncertainty is included, is denoted by EEV, and will be explained in more detail below.

Alternatively, the wait-and-see problem consists of solving a set of scenario problems, one for each possible outcome, one by one. Contrary to the recourse problem, no non-anticipativity constraints are included, meaning that the information process in the problem is completely disregarded. The expected value of the set of optimal solutions to the scenario problems is referred to as the *wait-and-see* (WS) solution, and represents the expected solution if all uncertainty is resolved. The set of scenario problems is often referred to as the WS model.

Finally, the value RP, also known as the *here and now solution*, denotes the optimal solution value to the recourse problem. Using the notation introduced, EVPI and VSS can be defined as follows:

$$EVPI = WS - RP \tag{3.3.1}$$

$$VSS = RP - EEV \tag{3.3.2}$$

EVPI, obtained by comparing the wait-and-see and here-and-now approaches, gives an indication of whether it is worth making an effort to reduce the uncertainty present in the problem. Calculating EVPI can therefore be useful for measuring the potential value of forecasting and decision support tools and the

information they provide (Guttormsen and Forsberg, 2004). VSS compares the expected value of using the expected value solution with the here-and-now solution, meaning that a small VSS implies a good approximation of the random variables by their expected values (Escudero et al., 2007). Figure 3.4 shows the relationship between the three mentioned solutions, and how they are used to define VSS and EVPI.

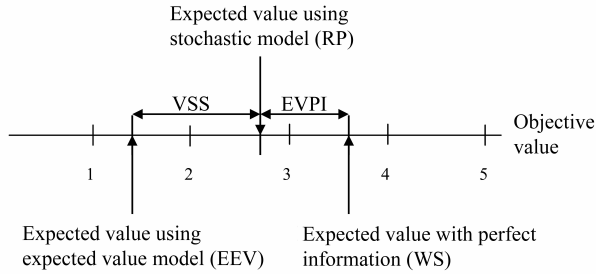


Figure 3.4: Relation between VSS and EVPI (Birge and Louveaux, 1997)

Figure 3.4 illustrates a general property that is valid for all recourse problems:

$$EEV \leq RP \leq WS \quad (3.3.3)$$

Two relationships must hold for the property above to be valid. $RP \geq EEV$ must be true, or RP is not the optimal solution to the recourse problem, since the expected value solution also is valid for the recourse problem, and therefore could have been chosen to obtain a better solution. $WS \geq RP$ is valid since the optimal solution to one outcome of the uncertain parameters is always better than or equal to the stochastic solution for the same outcome. Since EEV is the expectation over all outcomes, the inequality holds. For additional properties that are valid for recourse problems having specific structural properties, see Birge and Louveaux (1997).

So far, the value EEV has only been defined, without a further explanation of how this value is calculated. For two-stage models, EEV can be obtained by first solving the expected value problem, giving the EV solution. Second, the first-stage variables in the EV solution are used as the first-stage solution for each scenario in the WS model. Finally, the resulting scenario problems are solved, and EEV is given as the expectation over the set of solutions.

Due to the method in which the value EEV is calculated, there are several complications associated with applying VSS to multistage problems. The most apparent challenge is that for multistage problems, there is no obvious way of deciding which variables to fix in the WS model when calculating the value EEV . This, along with other complications, is dealt with in Escudero et al. (2007), where an

extension of VSS to multistage problems is presented. The parts of this study that are relevant to the model introduced in this next chapter will now be explained.

3.3.2 VSS in Multistage Models

Escudero et al. (2007) introduce the topic of VSS in multistage models by discussing which variables should be fixed in the WS model when calculating the value EEV . First, to emphasize that the existence of multiple stages has important implications, a trivial solution is suggested where variables are fixed in the WS model in the same manner as for two-stage models. That is, only the first-stage solutions are fixed, while the variables in the remaining stages are free to adapt to outcomes in the individual scenarios. This however, often results in a negative VSS, meaning that the WS model with fixed first-stage variables performs better than the recourse model. The reason for this is that non-anticipativity constraints are included in the recourse model, while in the WS model, each of the scenarios are solved as if they are independent. The example clearly illustrates that to successfully extend the notion of VSS to multistage problems, the value EEV must be redefined, and the non-anticipativity constraints must be considered. Escudero et al. (2007) suggest two approaches to how this can be done.

Approach A: The Value of the Stochastic Solution in t

In order to adapt the concept of EEV to a multistage setting, Escudero et al. (2007) begin by introducing the value EEV_t . EEV_t is fundamentally similar to the value EEV used for two-stage models, but differs in two ways. Firstly, instead of inserting the EV solution into the WS problem, EEV_t is calculated by inserting the EV solution into the recourse problem, RP, as given by equations 3.2.5 through 3.2.7. In this way, the non-anticipativity constraints are taken into account. Secondly, EEV_t has adopted the subscript t , denoting stage. These two modifications lead to EEV_t being defined as the optimal value of the RP model, where the decision variables until stage $t - 1$ are fixed at the optimal values obtained in the solution of the expected value problem (the EV solution).

$$EEV_t = \begin{cases} \text{RP model} \\ \text{s.t. } x_\omega^1 = \bar{x}^1, \omega \in \Omega \\ \dots \\ x_\omega^{t-1} = \bar{x}^{t-1}, \omega \in \Omega \end{cases} \quad (3.3.4)$$

Here, $\bar{x}^1, \bar{x}^2, \dots, \bar{x}^{t-1}$ are the optimal values obtained in the solution of the expected value problem. The set of scenarios, $\omega \in \Omega$, are defined as before.

In the same manner as for EEV_t , the measure VSS adopts a subscript t . VSS_t , the value of the stochastic solution in t for multistage models, is then defined as:

$$VSS_t = RP - EEV_t, \quad t = 1, \dots, T \quad (3.3.5)$$

VSS_t can be thought as a measure of the cost of ignoring uncertainty until stage t in the decision making process, that is, the performance of the approximation of the random variables by their expected values up to stage t .

Approach B: The Dynamic Value of the Stochastic Solution

The aim of the second approach presented by Escudero et al. (2007) is to derive a more realistic value for expected value solution, EV, by redefining the concept of the expected value problem. By dividing the scenario tree into sub-trees, a set of expected value (sub)problems can be solved, one for each sub-tree, hence the term dynamic. This allows for a more realistic utilization of the information process in the problem, which in turn increases the value of the information provided by the measure VSS. Before this method is explained in detail, some new notation must be introduced.

Since the decisions for all periods in a stage can be taken as soon as uncertainty is revealed in the first period, Escudero et al. (2007) adopt a simplified definition of the scenario tree where all periods in a stage are represented by a single node. These nodes are instead referred to as *scenario groups*, emphasizing the fact that a node in the scenario tree represents a group of scenarios when the recourse problem is represented as a set of subproblems (extensive form). G is used to denote the set of scenario groups g , where two scenarios belong to the same scenario group g in a given stage provided that they have the same realizations of the uncertain parameters up to that stage. G_t denotes the set of scenario groups in stage t , while the set of scenarios that belong to scenario group g are given by Ω_g .

Using the notation introduced, the basis for the approach can be described as defining an expected value problem for each scenario group, starting with the scenario group represented by the root node, and continuing down the tree, stage by stage. The expected value problem for a scenario group g is denoted by EV_g , which has the optimal solution Z_{EV}^g . The solution obtained by solving the expected value problems for the set of scenario groups in a given stage is referred to as the *dynamic solution of the expected value problem*.

This leads to the definition of $EDEV_t$, the expected result in t of using the dynamic solution to the expected value problem, as the expected value of the optimal solutions of the set of EV_g problems, where $g \in G_t$. Formally,

$$EDEV_t = \sum_{g \in \mathcal{G}_t} w^g Z_{EV}^g, \quad t = 1, \dots, T \quad (3.3.6)$$

where Z_{EV}^g is the optimal solution for the model EV_g and w^g represents the likelihood of the scenario group g , obtained as $w^g = \sum_{\omega \in \Omega_g} w^\omega$.

Finally, the dynamic value of the stochastic solution VSS^D , is defined as follows for all periods t in the last stage:

$$VSS_t^D = RP - EDEV_t \quad (3.3.7)$$

Chapter 4

Model Introduction

Before presenting the complete multistage stochastic model in the next chapter, some of the fundamentals in the model will now be explained. The aim of this chapter is to introduce the essence of the model and the setting in which it is used, without focusing on the mathematical notation. In this chapter and the next, the explanation of the model is done in a general manner, while chapters 6 and 7 demonstrate a more detailed reasoning. Section 4.1 introduces the model objective, section 4.2 explains the scope of the model, section 4.3 presents how the model is meant to be used, and the input given to the model is explained in section 4.4.

4.1 Model Objective

The aim of the model presented is to provide salmon producers with a tool that can aid them in making profitable decisions regarding harvesting and future sales. The tactical planning problem consists of when to harvest the salmon, and how to allocate sales between different contracts and sales in the spot market. Decisions must take into consideration the uncertainty in both future salmon prices and the growth of salmon, along with the many constraints that together define salmon aquaculture.

As opposed to the maximizing biomass output, which is the most commonly used approach in the industry, the objective of the model is to maximize the profits from sales. Maximizing profits bears great resemblance to maximizing biomass output since all profit comes from the sale of fish, but has the additional advantage of taking future price development into consideration. One of the reasons that maximizing biomass output has been the traditional approach is that until recently there was no efficient marketplace that could be used as a basis for estimating future prices, making planning of future sales more difficult. With the development of Fish Pool, along with increased access to market data

on supply and demand, salmon companies can adapt a more market oriented view on planning. This approach takes more information into account, thereby leading to more well informed decisions.

A third model objective alternative would be to minimize production risk, price risk, or both, given a chosen risk measure and certain sales or biomass constraints. But, based on the results of the empirical study done by Bergfjord (2009), an objective of minimizing risk would most likely not realistically portray the mindset of a salmon farmer. The results in the study show that even though salmon farmers themselves view salmon farming as riskier than other industries, they perceive themselves to be only moderately risk-averse, and more willing to take risk than farmers of land-based agriculture. Instead of including risk in the objective function, salmon farmers are assumed risk-neutral, and the model only takes risk into account through its stochastic nature, as discussed in section 3.3. Both the uncertainty in future prices and biomass development are included by letting the price and biomass variables be stochastic in the model.

4.2 Model Scope

The scope of the model is limited to only dealing with the activities that are directly associated with harvesting and sales. As opposed to a completely integrated model which would include all activities from breeding to final processing, only activities in the final stage of the seawater part of the value chain are modelled. The model takes as input the results of decisions made before smolt are released in seawater, and also assumes an optimal seawater production process. By this it is meant that seawater operations like feeding, keeping the fish healthy, and preventing escape are not included in the model. This limits the model to only consider decisions regarding what fish to sell, as well as when and how to sell it.

In each period, the producer is faced with four decisions; two concerning sales, and two concerning the fulfillment of delivery commitments. In regards to sales, the producer must decide which contracts to enter, and if a contract is to be entered, what amount to sell in the contract. This amount is restricted by an upper and lower bound, specified in the contract terms. Contracts are available for different fish classes, where classes are used as a method for classifying fish depending on a set of characteristics, for example size. The same applies to spot sales; both the fish class and amount of fish to be sold must be decided. Based on the spot sales made in a period and the contracts entered in previous periods, the producer must in each period decide how to supply the fish that is to be delivered. Delivery commitments can be fulfilled by deciding to harvest fish or by choosing to purchase fish in the spot market.

Limiting the scope of the model allows for only taking into account costs which are a direct consequence of the decisions made in the model. Depending on the timing and the method of sales, some costs may vary, while others are independent

of the decision variables and can be assumed constant. By disregarding all costs related to parts of the value chain that are not effected by sales decisions, the model is greatly simplified. Examples of direct costs that are not included in the model are all costs that occur before the smolt are released in seawater, as well as all costs associated with harvesting, transportation, maintenance, and overhead. It can be argued that all of the biomass must in any case be harvested, transported, and processed before delivery, so that omitting these costs will not affect the decisions made as long as the net present value implication of when these costs occur is disregarded. On the other hand, the cost of feeding the biomass from one period to the next is included, as this is a direct result of whether fish are harvested or not. The cost of feed makes up by far the largest share of the total costs, and even though some error inevitably is introduced by ignoring the remaining costs, some simplifications are necessary in order to make the model manageable. Due to the mentioned costs being omitted, the objective value of the profit maximization is only of comparative value, and must not be interpreted as the companies' net profit.

4.3 Using the Model

The model is meant to be used with a rolling horizon, as introduced in section 3.2.4. After implementing the first stage decisions, the model can be run again as soon as new information is received and additional decisions become available. For a salmon producer, new information about growth and price developments may imply adjustments to contract terms or the need for refinements in the harvesting plan, allowing for new decisions to be made. Since in reality this is a continuous process, new information will likely become available more often than what the model resolution can depict. The resolution in the model must therefore as accurately as possible portray this information and the decision process, and at the same time take into consideration that computational demand increases with increasing resolution. The amount of detail in the resolution is also limited by the processing times involved with harvesting and slaughtering. Section 6.1 gives an example of an actual planning problem modelled with a one month resolution.

In the same planning problem a 12 month planning horizon is used. A 12 month planning horizon has the benefit of including a complete cycle of the seasonal variations that are an important part of the industry. Using a longer planning horizon can prove difficult due to the unavailability of information. This is especially true for the sets of available contracts, since contracts in the industry normally have a maturity of 12 months or less (Asche, 2010). In addition, the quality of the model output depends on the quality of biomass and price estimates and realistic smolt release plans. The uncertainty introduced in these estimates as the time horizon is extended limits how far into the future it is reasonable to plan harvesting and sales. At the same time, it is necessary to allow for a sufficiently long time frame to avoid making sub-optimal short term decisions. Reducing the

planning horizon to much less than a year should therefore be avoided.

The model is intended to be used by companies that have a portfolio of locations large enough to create the need for an analytical approach to sales planning. For smaller companies where both the number of locations and number of contracts available are manageable, the value of the model is limited. Regardless of size, it is a requisite that the firm has forecasting capabilities at its disposal, either within the firm or provided by external forecasting services.

4.4 Model Input

The majority of the input data is a result of the discussion above regarding the scope of the model.

Biomass Development

A growth model specifies the growth rate for different fish classes, given certain conditions. By specifying the conditions in each region for each period in the model, the growth rates for each fish class in each region in each period are given. In addition to the growth model, biomass development is also affected by fish mortality and escape. These risks are modelled by including a survival rate for each fish class in each region in each period. Both the growth model and the survival rate are stochastic in the model, making the biomass development for periods in future stages uncertain.

Deterministic Parameters

Two deterministic parameters are given as input to the model; the processing ratios for each fish class and the cost per period of caring for each fish class. The costs associated with caring for the fish consist mainly of feeding costs, while the processing ratios give the weight per fish after gutting divided by the weight per fish before gutting. Including the processing ratios ensures that the weight lost due to gutting is taken into consideration when calculating the number of fish that it is necessary to slaughter in order to fulfill a delivery.

Region and Location Data

Data regarding location and region characteristics are given as input to the model. Although in reality each location consists of a set of net cages, production within each location is aggregated in the model. Each location is part of a region, and both locations and regions have MAB limits. Location MAB limits are specified for each period to allow for possible following restrictions (MAB is set to zero). Each region is characterized by a given upper and lower slaughtering limit and specific growth conditions that apply to all locations within the region. All location and region data is deterministic, given by strategic long term decisions.

Initial Biomass, Release Plan and Harvesting Periods

For each location, the number of fish in each fish class at the start of the first period must be given as input. Also given as input is a smolt release plan that specifies the number of fish per fish class that are to be released in a given location in a given period. The initial biomass and the release plan, combined with the growth model introduced above, are the basis for deciding in which periods harvesting is to be permitted in each location. This input is needed to ensure that valid end-of-horizon conditions are included in the model. This will be explained in detail in the next chapter. The initial biomass, the release plan, and the set of harvesting periods are all deterministic in the model.

Contract and Price Data

Data regarding price estimates for each fish class in each period and the set of contracts that the producer may enter in each period is given as input to the model. The set of contracts can be interpreted in two ways: either, the set can consist of contracts that actually are or will be available. Alternatively, the set can include a much larger variety of different contracts, so that the model solution advises the producer on which contracts it should pursue in negotiations with customers. In addition, contracts that are already entered and that are to be delivered during the planning period must also be given as input. Spot prices and the set of available contracts are stochastic, meaning that price development and available contracts for periods in future stages are uncertain.

Chapter 5

A Multistage Stochastic Model

A linear multistage stochastic model is now presented. The introduction of the model in the previous chapter along with the discussion of the situation faced by salmon farmers in chapter 2 provide the necessary background information for this chapter's explanation of the complete model. Section 5.1 explains how growth and decision variables are modelled in a manner that keeps the formulation linear, section 5.2 defines all sets, indexes, constants, and variables, while the objective function and the constraints are presented in section 5.3.

5.1 Modelling Growth and Decision Variables

Accurately modelling the situation faced by a salmon producer who seeks to optimize profits requires an overwhelming degree of detail. Therefore, assumptions and simplifications are essential in making the task feasible and developing a computationally tractable model. In addition, industry players are restrictive in terms of sharing information, especially regarding contract terms and conditions. The lack of information and the need for simplifications result in the following method of modelling growth and decision variables, presented along with an introduction of relevant notation.

Classification of Fish

The industry classifies fish using a set of characteristics such as age, weight, vaccines, feed type, quality, etc. In the model, weight is the only characteristic included, as this is by far the most important parameter, both for producers and for customers (Marine Harvest, 2011). For producers, knowing the weight of the fish is necessary in order to be able to decide the type and amount of feed, perform

growth estimates, and calculate biomass accounts. From a customer perspective, fish of different weights are suitable for different uses, and are thereby viewed as different products. To establish a connection between the way producers and customers deal with weight, fish are classified in two ways: as part of a fish class f in the set of fish classes \mathcal{F} and as part of a sales class j in the set of sales classes \mathcal{J} . The set of sales classes are defined by the market, while the set of fish classes represents a discretization of fish weight, which in reality of course is continuous.

Each fish class f is defined by a specific weight V_f , which is the weight of all fish in f at the beginning of a period. By in each period classifying fish as part of a fish class f with a defined weight, the task of modelling the weight of each fish is greatly simplified. In reality, V_f would be the mean weight of all fish within an interval defined by a lower boundary $\frac{V_f - V_{f-1}}{2}$ and an upper boundary $\frac{V_{f+1} - V_f}{2}$. In the model however, the variation in growth within a class, discussed in section 2.4, is ignored. All fish in a fish class f are assumed to grow at the same rate, making V_f the weight of all fish in fish class f .

A sales class j is a set consisting of one or more fish classes f . $\mathcal{F}(j)$ denotes the subset of fish classes belonging to sales class j . All sets j in \mathcal{J} are disjoint. For example, a fish may be in fish class f with $V_f = 3.75$ kilograms, where f is an element in the set j that consists of all fish classes with a V_f ranging from 3.00 to 5.00 kilograms.

Growth: Biomass Development

The basis for keeping track of biomass development in the model is the definition of the set \mathcal{F} of fish classes f as introduced above. In each period, each fish belongs to a fish class f , and as fish grow, they advance from one fish class to another. The number of individual fish in fish class f at the beginning of period t in location i in region r in scenario s is given by the decision variable m_{firs}^t . m_{firs}^t gives the number of fish at the beginning of a period, since this number may be reduced during a period due to mortality and escape. Mortality and escape are modelled by ϵ_{firs}^t , the survival rate for a fish in fish class f in region r in period t , in scenario s . Keeping track of the number of fish is in reality an integer-programming problem. To avoid the complexity involved with integer constraints, the variables m_{firs}^t are allowed to take on real values, a just simplification due to the magnitude of the number of fish.

The growth in kilograms in period t for fish in fish class f in region r in scenario s is given by the stochastic variable σ_{firs}^t . σ_{firs}^t is stochastic in order to reflect the uncertainty in salmon growth. The weight at the end of a period t for fish in fish class f in region r in scenario s is given by adding the growth during t , σ_{firs}^t , to the weight at the beginning of the period, V_f . The value $(V_f + \sigma_{firs}^t)$ will fall between two weights as defined by the set of fish classes.

$$V_{\underline{f}} \leq (V_f + \sigma_{f_{rs}}^t) \leq V_{\bar{f}} \quad (5.1.1)$$

Here, $V_{\underline{f}}$ is the weight of a fish in fish class \underline{f} , the fish class with defined weight closest to $(V_f + \sigma_{f_{rs}}^t)$ from below. $V_{\bar{f}}$ is the weight of a fish in fish class \bar{f} , the fish class with defined weight closest to $(V_f + \sigma_{f_{rs}}^t)$ from above.

As explained above, the model keeps track of the weight of each fish by in each period letting each fish belong to a fish class f with a defined weight V_f . Therefore, after each period, all fish must be distributed into new fish classes in a manner that accurately models the total biomass development during the last period. This is done by distributing the fish that at the end of a period t weigh $(V_f + \sigma_{f_{rs}}^t)$ between the two classes \underline{f} and \bar{f} , given by equation 5.1.1. The distribution is done based on how $(V_f + \sigma_{f_{rs}}^t)$ compares to $V_{\underline{f}}$ and $V_{\bar{f}}$, as shown in equation 5.1.2. $\delta_{r f \underline{f} s}^t$ is the share of fish class f that is distributed to fish class \underline{f} , and $\delta_{r f \bar{f} s}^t$ is the share of fish class f that is distributed to fish class \bar{f} .

$$\begin{aligned} \delta_{r f \bar{f} s}^t &= \frac{V_{\bar{f}} - (V_f + \sigma_{f_{rs}}^t)}{V_{\bar{f}} - V_{\underline{f}}} \\ \delta_{r f \underline{f} s}^t &= \frac{(V_f + \sigma_{f_{rs}}^t) - V_{\underline{f}}}{V_{\bar{f}} - V_{\underline{f}}} \end{aligned} \quad (5.1.2)$$

Equations 5.1.1 and 5.1.2 are part of the data preprocessing, not the model itself. They are included here to ease the understanding of how $\delta_{f_{rs}}^t$ is calculated.

Accurately modelling biomass development requires a combination of sufficiently large time resolution and sufficiently detailed partitioning of fish weight. The reason for this is that if the growth of fish in fish class f during period t in region r , $\sigma_{f_{rs}}^t$, is such that $\underline{f} = f$ for $\delta_{r f \underline{f} s}^t$, a share of the fish will remain in f . If this repeats for every period, some of the fish will never grow to become larger than f , a situation which clearly is not acceptable. To prevent this from occurring, \underline{f} must in most conditions be larger than f , so that all fish advance to a new class. This is ensured by having a sufficiently long period duration combined with a small enough difference in weight from one fish class to the next. A longer period duration implies a larger $\sigma_{f_{rs}}^t$, while the more detailed the partitioning, the less the fish need to grow in order to advance to the next class.

Decision Making

All decisions in a period t are taken at the start of a period. The number of fish in each fish class f is updated based on the growth development and survival rate in period $t - 1$. Smolt that are to be released in period t are released, new

future contracts are entered, and spot sales are executed. The fish that are to be delivered in period t are either harvested and slaughtered or purchased in the spot market. All of these activities are done without a particular order, meaning that the decisions can be regarded as being taken simultaneously.

When a decision is made regarding harvesting in month t , it is assumed that the effects of harvesting occur immediately. This implies that the processing times normally associated with sorting, harvesting, transporting, starving, slaughtering, and delivering the fish, are neglected, and the process from sorting to delivery can be regarded as instantaneous in month t . Therefore, the terms harvesting and slaughtering are from now on used interchangeably. The limitations and costs that apply to sorting capabilities and transportation with well boats are omitted from the model, as explained in section 4.2.

Contracts

A contract is defined as an agreement to deliver a certain number of kilograms of a specified sales class at a given point in time in the future. α_{jdc}^t describes the characteristics of a contract as the price per kilogram for contract c of sales class j with delivery in period d , entered in period t , in scenario s . The decision variable x_{jdc}^t has the same indexes, and denotes the number of kilograms sold in contract c at price α_{jdc}^t . For each contract, a maximum and a minimum number of kilograms are given by μ_{jdc}^t and ι_{jdc}^t , respectively. The uncertainty in future prices and demand are taken into account by defining α_{jdc}^t , μ_{jdc}^t , and ι_{jdc}^t as stochastic variables.

An alternative method of modelling contracts is fixing the size of each contract and using a single binary variable to denote whether a contract is entered or not. The formulation used in the model is preferable because it avoids the additional complexity of dealing with binary variables. As with m_{firs}^t , the variables x_{jdc}^t are also allowed to take on real values. Contracts that are entered before period $t = 1$ and that are to be delivered during the planning period are modelled by including a deterministic parameter α_{jdc}^t for each contract, and by forcing x_{jdc}^t to take on the same given input value for all scenarios s .

Spot Sales and Purchases

Spot sales and purchases are defined in the same manner as for contracts. β_{js}^t and γ_{js}^t denote respectively the spot price in period t for sales and purchases of sales class j in scenario s . z_{js}^t denotes the number of kilograms of sales class j sold spot at price γ_{js}^t in period t in scenario s , while y_{js}^t denotes the number of kilograms of sales class j purchased spot at price β_{js}^t in period t in scenario s .

The prices for spot sales and spot purchases are modelled separately. This is done in order to model the costs that may incur as a result of the uncertainty in biomass development, as discussed in section 2.4. To include the costs associated with a producer being forced to purchase fish in the spot market, the spot purchase

price is modelled above that of the spot sales price. The possible costs involved with having to sell fish in the spot market are modelled by including a premium in the contract prices. The option to sell and purchase fish in the spot market also have the purpose of ensuring feasibility in the model. Spot sales allow for selling fish when necessary in order to satisfy biomass restrictions, while allowing for spot purchases in all periods guarantees the delivery of all contracts.

5.2 Formal Definitions

The following section formally defines all sets, indexes, constants, and variables that are included in the model. Sets are denoted by capital, calligraphic letters with corresponding indexes. Deterministic data (constants) are capitalized, decision variables are denoted in lower case Latin letters, while Greek letters are used to denote stochastic data. Quantities refer to a number of fish, while amounts describe weight or biomass, and is given in kilograms. Prices are per kilogram.

Sets

\mathcal{F}	Set of fish classes
\mathcal{J}	Set of sales classes
$\mathcal{F}(j)$	Set of fish classes f in sales class j , $\mathcal{F}(j) \subset \mathcal{F}$
$\mathcal{I}(r)$	Set of locations in region r
\mathcal{R}	Set of regions
\mathcal{S}	Set of scenarios
\mathcal{T}	Set of periods
$\mathcal{T}(i)$	Set of periods when harvesting in location i is <i>not</i> permitted, $\mathcal{T}(i) \subset \mathcal{T}$
$\mathcal{C}(j, d, t)$	Set of contracts available at time t for sales class j with delivery in period d
\mathcal{D}	Set of contract delivery periods, $\mathcal{D} \subset \mathcal{T}$
\mathcal{N}	Set of event nodes in the scenario tree

Indexes

f, \hat{f}	Index for fish class
j	Index for sales class
i	Index for location
r	Index for region
s	Index for scenario
t, \hat{t}	Index for period
c	Index for contract
d	Index for contract delivery period
n	Index for event node

Deterministic Data

K_f	Cost per period per kilogram of caring for fish class f
N_{fir}^t	Number of smolt of fish class f released in location i in region r in period t
m_{firs}^1	Initial number of fish in fish class f in location i in region r , equal for all scenarios s
V_f	Weight in kilograms of a fish in fish class f at the beginning of a period
Q_f	Processing ratio (loss factor due to gutting) for fish class f
MAB_i^t	MAB in kilograms for location i in period t
MAB_r	MAB in kilograms for region r
\bar{S}_r	Maximum slaughtering quantity per period in region r
\underline{S}_r	Minimum slaughtering quantity per period in region r

Stochastic Data

α_{jdc}^t	Price per kilogram for contract c of sales class j with delivery in period d , entered in period t , in scenario s
γ_{js}^t	Price per kilogram for spot sales of sales class j in period t , in scenario s
β_{js}^t	Price per kilogram for spot purchases of sales class j in period t , in scenario s
μ_{jdc}^t	Maximum number of kilograms in contract c of sales class j with delivery in period d , entered in period t , in scenario s
l_{jdc}^t	Minimum number of kilograms in contract c of sales class j with delivery in period d , entered in period t , in scenario s
c_{frr}^t	The survival rate for a fish in fish class f in region r in period t , in scenario s
δ_{ffrs}^t	The share of fish class \hat{f} that has grown to become part of fish class f due to the growth in period t in region r , in scenario s
σ_{frr}^t	The growth in kilograms in period t for fish class f in region r , in scenario s
ρ_s	The probability of scenario s

Decision Variables

x_{jdc}^t	Number of kilograms of sales class j sold in contract c at price α_{jdc}^t with delivery in period d , entered in period t , in scenario s
z_{js}^t	Number of kilograms of sales class j sold spot at price γ_{js}^t in period t , in scenario s
y_{js}^t	Number of kilograms of sales class j purchased spot at price β_{js}^t in period t , in scenario s
w_{firs}^t	Number of fish in fish class f harvested from location i in region r in period t , in scenario s
m_{firs}^t	Number of fish in fish class f in location i in region r in the beginning of period t after harvesting and release, in scenario s

5.3 Model Formulation

The model will now be explained in detail, starting with the objective function and continuing with the different constraints. The model is formulated as a deterministic equivalent problem in extensive form, as explained in detail in section 3.2. This means that the complete problem consists of a set of subproblems, one subproblem for each scenario, with the problem structure enforced by a set of non-anticipativity constraints.

5.3.1 Objective Function

The model maximizes the expected income from contract sales and sales in the spot market, minus the cost of spot purchases and cost of caring for the biomass.

$$\begin{aligned} \text{Max}_x \sum_{s \in \mathcal{S}} \rho_s & \left(\sum_{t \in \mathcal{T}} \left(\sum_{j \in \mathcal{J}} \sum_{d \in \mathcal{D}} \sum_{c \in \mathcal{C}(j,d,t)} \alpha_{jdc}^t x_{jdc}^t + \sum_{j \in \mathcal{J}} \gamma_{js}^t z_{js}^t \right. \right. \\ & \left. \left. - \sum_{j \in \mathcal{J}} \beta_{js}^t y_{js}^t - \sum_{r \in \mathcal{R}} \sum_{i \in \mathcal{I}(r)} \sum_{f \in \mathcal{F}} K_f V_f m_{firs}^t \right) \right) \end{aligned} \quad (5.3.1)$$

Here, ρ_s is the probability of scenario s . For each scenario s , in each period t , four terms are added. In the first term, x_{jdc}^t denotes contract sales, the number of kilograms of sales class j with delivery in period d , sold in contract c at price α_{jdc}^t . In the second term, income from spot sales, z_{js}^t is the number of kilograms of sales class j sold spot at price γ_{js}^t . The third term represents the costs of spot purchasing, where y_{js}^t is the number of kilograms of sales class j purchased in the spot market at price β_{js}^t . The fourth and final term is the total cost of caring for the biomass in period t , where K_f is the cost per period per kilogram of caring for fish class f , V_f is the weight of a fish in fish class f at the beginning of a period, and m_{firs}^t is the number of fish in fish class f in location i in region r at the beginning of a period t after harvesting and release.

Regarding the first term, since all contracts entered into will be delivered (guaranteed by spot purchasing) and the time value of money is disregarded, the income from a contract sold can be included at the time of entering (as opposed to at the time of delivery) without affecting the results.

5.3.2 Constraints

The following constraints are included in the model. Since all constraints apply to all scenarios s , the index s for scenario will be omitted in some of the explanations in the remainder of this section for ease of reading.

Initial Conditions

In period $t = 1$, the only decision to be made is regarding entering contracts. It is assumed that the decisions regarding spot sales, spot purchases, and harvesting have already been made and are reflected in the initial biomass, m_{firs}^1 . m_{firs}^1 is the initial number of fish in fish class f in location i in region r in scenario s , which is given as an input to the model and is equal for all scenarios s .

$$z_{js}^1 = 0, \quad j \in \mathcal{J}, \quad s \in \mathcal{S} \quad (5.3.2)$$

$$y_{js}^1 = 0, \quad j \in \mathcal{J}, \quad s \in \mathcal{S} \quad (5.3.3)$$

$$w_{firs}^1 = 0, \quad r \in \mathcal{R}, \quad i \in \mathcal{I}(r), \quad f \in \mathcal{F}, \quad s \in \mathcal{S} \quad (5.3.4)$$

Here, z_{js}^t is the number of kilograms of sales class j that is sold spot in period $t = 1$, y_{js}^t is the number of kilograms of sales class j that is purchased spot in period $t = 1$, and w_{firs}^t is the number of fish in fish class f harvested from location i in region r in period $t = 1$.

Biomass Development

The following constraint keeps track of the development in biomass from one period to the next. This is done by keeping track of which fish class f each fish belongs to, and how growth affects the advancement of fish from one fish class f to another. m_{firs}^t , the number of fish in a fish class f in a location i in a region r at that beginning of a period t , is determined by the following three elements:

1. The number of fish that are in fish class f after the biomass development due to growth and survival rate during period $t - 1$.
2. The number of fish released in fish class f in period t .
3. The number of fish harvested in fish class f in period t .

Each of these elements constitute a term in the equation below. The constraint does not apply to period $t = 1$ since m_{firs}^1 , the initial biomass for fish class f in location i in region r , is given as an input to the model, and there is no release or harvesting in the first period.

$$m_{firs}^t = \sum_{\hat{f} \leq f} (\delta_{\hat{f}firs}^{t-1} m_{\hat{r}\hat{i}\hat{f}s}^{t-1} \epsilon_{\hat{f}firs}^{t-1}) + N_{fir}^t - w_{firs}^t, \quad (5.3.5)$$

$$t \in \mathcal{T} \setminus \{1\}, \quad r \in \mathcal{R}, \quad i \in \mathcal{I}(r), \quad f \in \mathcal{F}, \quad s \in \mathcal{S}$$

$\delta_{\hat{f}f_{rs}}^{t-1}$ is the share of fish class \hat{f} that during period $t - 1$ has grown to become part of fish class f in region r . This share is multiplied by $m_{rif_s}^{t-1}$, the number of fish in fish class \hat{f} in location i in region r at the beginning of period $t - 1$, and $\epsilon_{f_{rs}}^{t-1}$, the survival rate for a fish in fish class \hat{f} in region r in period $t - 1$. When calculating the number of fish in a fish class f , the sum over all fish classes smaller than f add to f all of the fish that have grown to become part of f during the last period. Also, there is the possibility that fish may remain in the same fish class from one period to the next. By including the fish class f itself in the sum, $\delta_{ff_{rs}}^t$ ensures that the fish that remain in f are also included. For a detailed explanation of the calculation of $\delta_{\hat{f}f_{rs}}^{t-1}$, see section 5.1.

N_{fir}^t is the number of smolt in fish class f released in location i in region r in period t . w_{firs}^t is the number of fish in fish class f harvested from location i in region r in period t .

Contract Size

The size of each contract is defined by a minimum number of kilograms which must be delivered and a maximum number of kilograms which the delivery cannot exceed. If the minimum amount equals the maximum amount, the contract of course specifies an exact amount.

$$l_{jdc_s}^t \leq x_{jdc_s}^t \leq \mu_{jdc_s}^t, \quad t \in \mathcal{T}, j \in \mathcal{J}, d \in \mathcal{D}, c \in \mathcal{C}(j, d, t), s \in \mathcal{S} \quad (5.3.6)$$

Here, $x_{jdc_s}^t$ is the number of kilograms of sales class j sold in contract c with delivery in period d , entered in period t , restricted by the minimum amount $l_{jdc_s}^t$ and the maximum amount $\mu_{jdc_s}^t$.

Sales and Biomass

The number of kilograms of fish in each sales class j that is to be delivered in a period t must be equal to the number of kilograms harvested (adjusted for gutting) plus the number of kilograms supplied by spot purchasing in the same period. This is enforced through the constraint below, which has several important implications in the model:

1. The constraint provides a connection between sales classes j and fish classes f . To supply fish for delivery in sales class j , fish may be harvested from all fish classes f which are elements in j .
2. Fish in a fish class f are often available in several locations at the same time. The choice of from which location fish are to be harvested is enforced by this constraint.
3. Fish are purchased in the spot market when the amount to be delivered cannot be supplied by harvesting, or when it is profitable not to harvest

available fish.

Each of these elements constitute a term in the equation below. The constraint does not apply to period $t = 1$ since delivery for this period is already completed.

$$\sum_{\hat{t}=1}^{\hat{t}<t} \sum_{c \in \mathcal{C}(j,d,t)} x_{j\hat{t}cs}^{\hat{t}} + z_{js}^t = \sum_{r \in \mathcal{R}} \sum_{i \in \mathcal{I}(r)} \sum_{f \in \mathcal{F}(j)} Q_f V_f w_{firs}^t + y_{js}^t, \quad (5.3.7)$$

$t \in \mathcal{T} \setminus \{1\}, j \in \mathcal{J}, s \in \mathcal{S}$

Here, $x_{j\hat{t}cs}^{\hat{t}}$ is the number of kilograms of sales class j sold in contract c , entered in period \hat{t} with delivery in period t . z_{js}^t is the number of kilograms of sales class j that is sold spot in period t , y_{js}^t is the number of kilograms of sales class j that is purchased spot in period t , and w_{firs}^t is the number of fish in fish class f harvested from location i in region r in period t . Q_f is the ratio of the weight of a fish in fish class f after gutting to the weight of the fish before gutting, included so that the loss due to gutting is taken into consideration. V_f is the weight of a fish in fish class f at the beginning of a period.

Slaughtering Capacity

The following constraint ensures that the producer does not sell more or less than what the slaughtering facilities can manage. Slaughtering capacity within a region is aggregated, and it is assumed that a location in region r must use slaughtering capacity in region r . The constraint does not apply to period $t = 1$ since delivery for this period is already completed.

$$\underline{S}_r \leq \sum_{i \in \mathcal{I}(r)} \sum_{f \in \mathcal{F}} w_{firs}^t \leq \overline{S}_r, \quad t \in \mathcal{T} \setminus \{1\}, r \in \mathcal{R}, s \in \mathcal{S} \quad (5.3.8)$$

Here, w_{firs}^t is the number of fish in fish class f harvested from location i in region r in period t . \underline{S}_r and \overline{S}_r are the minimum and maximum slaughtering quantities in region r , respectively.

Maximum Allowable Biomass and Fallowing Periods

To ensure that producers comply with the set biomass limits, it is necessary to restrict the amount of biomass in all periods t , both for individual locations i as well as for regions r as a whole. The restrictions must be enforced when the biomass is at its maximum, and since growth is strictly positive, the maximum biomass is always achieved at the end of a period. The constraints do not apply to period $t = 1$ since decisions affecting the biomass in the first period are already taken.

The first constraint below restricts the biomass in each location, while the second restricts the biomass in each region. In addition to enforcing the set biomass limits, the location restriction also ensures that fallowing is performed when necessary by setting MAB_i^t equal to zero. Fallowing periods can either be due to lice zone regulations that cover multiple locations, or due to fallowing requirements in a single location ahead of a planned release.

$$\sum_{f \in \mathcal{F}} (V_f m_{firs}^t + \sigma_{firs}^t m_{firs}^t) \epsilon_{firs}^t \leq MAB_i^t, \quad (5.3.9)$$

$$t \in \mathcal{T} \setminus \{1\}, r \in \mathcal{R}, i \in \mathcal{I}(r), s \in \mathcal{S}$$

$$\sum_{i \in \mathcal{I}(r)} \sum_{f \in \mathcal{F}} (V_f m_{firs}^t + \sigma_{firs}^t m_{firs}^t) \epsilon_{firs}^t \leq MAB_r, \quad (5.3.10)$$

$$t \in \mathcal{T} \setminus \{1\}, r \in \mathcal{R}, s \in \mathcal{S}$$

Here, the first term in each equation represents the biomass at the beginning of period t , while the second term is the biomass that results from growth in period t . Both terms are multiplied by ϵ_{firs}^t , the survival rate for a fish in fish class f in region r . V_f is the weight in kilograms of a fish in fish class f at the beginning of a period, and m_{firs}^t is the number of fish in fish class f in location i in region r at the beginning of period t , after harvesting and release. σ_{firs} is the growth in period t for a fish in fish class f in region r . MAB_i^t is the maximum allowable biomass for location i in period t , and MAB_r is the maximum allowable biomass for region r .

Harvesting Periods

End-of-horizon conditions are modelled by using problem specific business rules that restrict when harvesting is permitted in each location. Depending on the length of the planning period, some fish may not be ready for harvesting until after the final period. For example, with a planning period of twelve months, this would apply to both fish that are released just before and during the planning period. Without specifying in which periods harvesting is allowed for each location, sub-optimal harvesting decisions will be made for the final periods in order to maximize the planning period profit. The reason for choosing to restrict harvesting by using harvesting periods instead of restricting which fish classes that can be harvested is that the planned harvesting weight often varies between locations. And, since locations generally only contain fish within a limited size range, restricting all fish within a location is valid since they are all normally harvested approximately at the same time.

$$w_{firs}^t = 0, \quad r \in \mathcal{R}, i \in \mathcal{I}(r), f \in \mathcal{F}, t \in \mathcal{T}(i), s \in \mathcal{S} \quad (5.3.11)$$

The constraint ensures that w_{firs}^t , the number of fish in fish class f harvested from location i in region r in period t , is zero for all periods in $\mathcal{T}(i)$, the set of periods when harvesting in location i is not allowed.

Non-Anticipativity

The structure of the scenario tree and the relationship between stages, periods and scenarios is enforced by non-anticipativity constraints. As discussed in section 3.2, each stage is a subset of the set of periods, defined in accordance with the resolution of uncertainty in the decision problem. The non-anticipativity constraints force decisions in different scenarios s to be equal in a manner that is consistent with the information available in each period t . The scenarios passing through node n are given by $S(n)$, while $T(n)$ is the time period of node n (Rockafellar and Wets, 1991).

$$\begin{aligned} & \frac{1}{|S(n)|} \sum_{s' \in S(n)} (x_{jdc s'}^t, y_{j s'}^t, z_{j s'}^t, w_{firs'}^t, m_{firs'}^t) \\ & = (x_{jdc s}^t, y_{j s}^t, z_{j s}^t, w_{firs}^t, m_{firs}^t), \\ & j \in \mathcal{J}, d \in \mathcal{D}, c \in \mathcal{C}(j, d, t), r \in \mathcal{R}, i \in \mathcal{I}(r), \\ & f \in \mathcal{F}, n \in \mathcal{N}, s \in \mathcal{S}(n), t \in \mathcal{T}(n) \end{aligned} \quad (5.3.12)$$

5.3.3 The Complete Model

Finally, the complete model is presented, along with non-negativity constraints on all decision variables, where the indexes are omitted.

$$\begin{aligned} & \text{Max}_x \sum_{s \in \mathcal{S}} \rho_s \left(\sum_{t \in \mathcal{T}} \left(\sum_{j \in \mathcal{J}} \sum_{d \in \mathcal{D}} \sum_{c \in \mathcal{C}(j, d, t)} \alpha_{jdc s}^t x_{jdc s}^t + \sum_{j \in \mathcal{J}} \gamma_{j s}^t z_{j s}^t \right) \right. \\ & \left. - \sum_{j \in \mathcal{J}} \beta_{j s}^t y_{j s}^t - \sum_{r \in \mathcal{R}} \sum_{i \in \mathcal{I}(r)} \sum_{f \in \mathcal{F}} K_f V_f m_{firs}^t \right) \end{aligned}$$

s.t.

$$z_{j s}^1 = 0, \quad j \in \mathcal{J}, s \in \mathcal{S}$$

$$y_{j s}^1 = 0, \quad j \in \mathcal{J}, s \in \mathcal{S}$$

$$w_{firs}^1 = 0, \quad r \in \mathcal{R}, i \in \mathcal{I}(r), f \in \mathcal{F}, s \in \mathcal{S}$$

$$m_{firs}^t = \sum_{\hat{f} \leq f} (\delta_{\hat{f}f}^{t-1} m_{r\hat{f}s}^{t-1} \epsilon_{\hat{f}rs}^{t-1}) + N_{fir}^t - w_{firs}^t,$$

$$t \in \mathcal{T} \setminus \{1\}, r \in \mathcal{R}, i \in \mathcal{I}(r), f \in \mathcal{F}, s \in \mathcal{S}$$

$$l_{jdc}^t \leq x_{jdc}^t \leq \mu_{jdc}^t, \quad t \in \mathcal{T}, j \in \mathcal{J}, d \in \mathcal{D}, c \in \mathcal{C}(j, d, t), s \in \mathcal{S}$$

$$\sum_{\hat{t}=1}^{\hat{t}<t} \sum_{c \in \mathcal{C}(j, d, \hat{t})} x_{jtc}^{\hat{t}} + z_{js}^t = \sum_{r \in \mathcal{R}} \sum_{i \in \mathcal{I}(r)} \sum_{f \in \mathcal{F}(j)} Q_f V_f w_{firs}^t + y_{js}^t,$$

$$t \in \mathcal{T} \setminus \{1\}, j \in \mathcal{J}, s \in \mathcal{S}$$

$$\underline{S}_r \leq \sum_{i \in \mathcal{I}(r)} \sum_{f \in \mathcal{F}} w_{firs}^t \leq \overline{S}_r, \quad t \in \mathcal{T} \setminus \{1\}, r \in \mathcal{R}, s \in \mathcal{S}$$

$$\sum_{f \in \mathcal{F}} (V_f m_{firs}^t + \sigma_{firs}^t m_{firs}^t) \epsilon_{firs}^t \leq MAB_i^t,$$

$$t \in \mathcal{T} \setminus \{1\}, r \in \mathcal{R}, i \in \mathcal{I}(r), s \in \mathcal{S}$$

$$\sum_{i \in \mathcal{I}(r)} \sum_{f \in \mathcal{F}} (V_f m_{firs}^t + \sigma_{firs}^t m_{firs}^t) \epsilon_{firs}^t \leq MAB_r,$$

$$t \in \mathcal{T} \setminus \{1\}, r \in \mathcal{R}, s \in \mathcal{S}$$

$$w_{firs}^t = 0, \quad r \in \mathcal{R}, i \in \mathcal{I}(r), f \in \mathcal{F}, t \in \mathcal{T}(i), s \in \mathcal{S}$$

$$\frac{1}{|S(n)|} \sum_{s' \in S(n)} (x_{jdc}^t, y_{js'}^t, z_{js'}^t, w_{firs'}^t, m_{firs'}^t)$$

$$= (x_{jdc}^t, y_{js}^t, z_{js}^t, w_{firs}^t, m_{firs}^t),$$

$$j \in \mathcal{J}, d \in \mathcal{D}, c \in \mathcal{C}(j, d, t), r \in \mathcal{R}, i \in \mathcal{I}(r),$$

$$f \in \mathcal{F}, n \in \mathcal{N}, s \in \mathcal{S}(n), t \in \mathcal{T}(n)$$

$$x, z, y, w, m, \geq 0$$

Chapter 6

Model Application

The model presented in the previous chapter is implemented in three different versions, all of which are applied to Marine Harvest Region Mid. Section 6.1 introduces the three models and details regarding the implementations. Section 6.2 presents Marine Harvest Region Mid and the rest of the input data. Finally, sections 6.3 and 6.4 describe in further detail the two-stage and multistage models, respectively. The purpose of implementing and applying the models is not to provide conclusive results as to when salmon should be harvested or how sales should be allocated. Rather, this chapter and the next intend to only demonstrate how the model presented in the previous chapter can be used to solve a realistic salmon sales planning problem.

6.1 The Models

Three versions of the multistage stochastic model presented in the previous chapter are now proposed. By comparing the results of applying these three models to a realistic salmon sales planning problem, the model in the previous chapter can be evaluated, and the implications of the uncertainty present in salmon farming further understood. The three models are a deterministic model that ignores the uncertainty in both price and salmon growth, and two stochastic models, a two-stage and a multistage, where both include the uncertainty in biomass development, while ignoring price stochasticity. In the *deterministic* (DET) model, expected values are used for the uncertain parameters, while in the *two-stage stochastic* (TS) and the *multistage stochastic* (MS) models, uncertainty is resolved during the planning period. The original model presented in the previous chapter will from now on be referred to as the *general* (GEN) model. Details regarding how the models are implemented and the simplifications made will now be discussed in further detail.

6.1.1 Simplifications

An important part of the implementation is the set of simplifications made from the GEN model, the most important one being that price uncertainty is omitted in the two stochastic models. This is necessary in order to make the stochastic models computationally tractable for a personal computer. The memory in the computer used for running the models can only handle a limited number of variables, and hence scenarios. If the uncertainty in biomass development is to be modelled with a sufficient degree of detail, it is not possible to also include price uncertainty in the current models without additional computational power. Ideas for a future extension of the models including price uncertainty are presented in chapter 8.

An additional simplification is that biomass development is assumed to depend only on seawater temperature, while disease, escape, and mortality are ignored. Since approximately 80% of fish mortality occurs immediately after release, it can be argued that this source of uncertainty can be partly dealt with through release planning, which is taken as input to the models. Disease outbreaks and escape are difficult to model realistically because of the randomness involved, and are left as a possible future extension to the models.

6.1.2 Implementation of the Models

All three models are written in Mosel and implemented in Xpress-IVE, version 1.21.02, 64 bit. Xpress-IVE is the interface or editor in the Xpress Product Suite by FICO. It allows for graphical displays of the solution and run-time statistics. Mosel is a compiled modelling and programming language specifically designed for rapid modelling of optimization problems. The models are solved by the Xpress Optimizer, which solves linear, continuous, integer, and mixed-integer programs using robust optimization algorithms (FICO, 2010). The Xpress product suite is presented in figure 6.1.

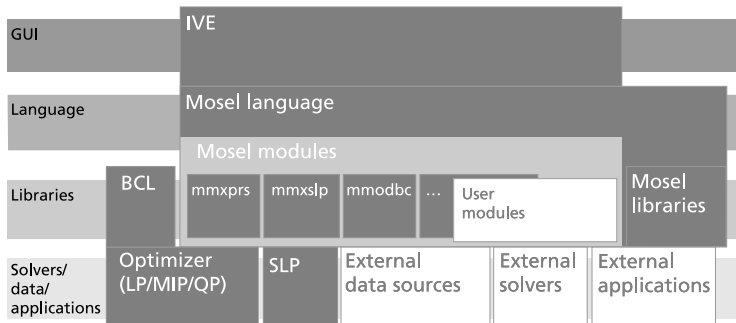


Figure 6.1: Xpress Product Suite (FICO, 2010)

In Xpress, the MMODBC module included in the ODBC interface has been used to exchange data with Microsoft Excel. All input data is read from an input spreadsheet, and the results are written to an output spreadsheet. A considerable effort has been made to structure both the input and output spreadsheets in a clear and understandable manner to ensure that the models can be used without detailed knowledge of modelling or optimization. All calculations are done using functions in Xpress, thus eliminating the need for manual pre-solve calculations of the input. In other words, after entering the input data specified in section 4.4 into the input spreadsheet, the Xpress models can be run, and relevant output is written to the output spreadsheet where a simple graphical analysis of the solution is presented. The algebraic notation in the GEN model is formulated in a manner that allows for an almost direct implementation in Xpress. This makes the GEN model a good resource for understanding the details of the implementation, thereby allowing for easy alteration or extensions.

All three models are run on a Dell XPS M1330 with an Intel Core 2 Duo T7500 2.2 GHz processor, 4GB RAM memory, and Microsoft Windows 7 Professional 64-bit operating system.

6.2 Data Set

6.2.1 Marine Harvest Region Mid

Marine Harvest is the world's largest producer of farmed salmon with approximately one fifth of the global production. Marine Harvest Norway has farming and processing activities along most of the Norwegian coastline, and is divided into four main regions: Region North, Region Mid, Region West and Region South. The models are applied to Marine Harvest Region Mid, which spans from Averøy in the south to Fosnes in the north, and employs approximately 230 people. The region consists of 39 seawater locations, five freshwater locations, three seawater broodstock facilities, one spawning facility, and one processing factory (Marine Harvest, 2010b). Only the seawater locations and the processing factory are relevant to the models, and from here on location refers to seawater location.

Since the regions set by Marine Harvest are not consistent with those of the Directorate of Fisheries, Region Mid actually consists of two regions in regards to regional MAB limits. The 6 most southerly locations belong to the region Møre og Romsdal, where the regional MAB limit is 7800 tons. The remaining 33 locations belong to the region Trøndelag, which has a MAB limit of 32760 tons. The MAB limits for the individual locations in the two regions range from 1520 to 7020 tons. Locations in both regions use the same slaughtering facility located in Hitra, where the slaughtering capacity is 70 000 fish per day, or a little more than 2 million fish per month. A synthetic but realistic lower slaughtering limit of 500 000 fish per month is included in the implementation. Recently,

restrictions regarding common fallowing for all locations in a given area have been established in parts of Marine Harvest Region Mid. These restrictions are not included in the implementation, as common fallowing is only enforced in certain parts of the country. The locations that belong to Marine Harvest Region Mid are summarized in table 6.1.

6.2.2 Resolution and Planning Horizon

The models have a planning horizon of one year, thereby including a complete cycle of the seasonal variations in temperature. The resolution in the models is one month, with the first period being January. The planning horizon and the resolution are chosen based on the forecasting capabilities available to Marine Harvest (Marine Harvest, 2011), and the discussion in section 4.3 regarding the factors limiting the amount of detail in the resolution and the length of the planning horizon. Contract sales are simplified by only including contracts that have delivery within the planning period. This simplification is possible because of how the end-of-horizon conditions are enforced, but limits the assortment of contracts in later periods, and implies no contract sales in period 12.

6.2.3 Biomass Development

Biomass development in the implementation is modelled by classifying fish using a total of 82 fish classes f . The implemented growth model is based on a model provided by Skretting, the world's leading producer of fish feed. The original Skretting growth model gives the daily growth in percent for different sized Atlantic salmon given temperatures ranging from 1 to 20 degrees Celsius. It is based on results from their R_{max} database that divides fish into 34 sizes. This resolution has been expanded to the 82 classes in \mathcal{F} in the implemented model, with the growth rates of new fish classes calculated based on weighted growth rates of the size above and below each new fish class. The temperature resolution in the implemented growth model has also been increased, from 1 degree Celsius in the Skretting model to 0.5 degrees Celsius in the implemented growth model. The new temperature steps have been calculated in the same way as for the new fish classes. All growth rates have been converted to kilograms per month. An excerpt of the implemented growth model and the set of fish classes \mathcal{F} is shown in table 6.2. Note that the final fish class, $f = 82$, has a growth rate of zero. This is done to ensure feasibility in the models, and has no effect on the decisions taken since no fish grow to become this size, as explained in more detail below.

6.2.4 Scenarios

With price assumed deterministic and growth depending solely on temperature, the growth scenarios are generated by varying the temperature input. The scenarios in the models are given by temperature data from Marine Harvest Region

Table 6.1: Locations in Marine Harvest Region Mid (Directorate of Fisheries, 2010)

<i>Location</i>	<i>Municipality</i>	<i>Location MAB</i>	<i>Region</i>
Bremsessvaet	Smøla	5460 tons	Møre og Romsdal
Brettingen	Smøla	5460 tons	Møre og Romsdal
Leite	Averøy	3900 tons	Møre og Romsdal
Kornstad	Averøy	3120 tons	Møre og Romsdal
Rokset	Averøy	3120 tons	Møre og Romsdal
Storvikja	Aure	3120 tons	Møre og Romsdal
Tennøya	Frøya	3900 tons	Trøndelag
Mannbruholmen	Frøya	7020 tons	Trøndelag
Grøttingsøy	Frøya	5460 tons	Trøndelag
Slettholmene	Frøya	3120 tons	Trøndelag
Langskjæra	Frøya	5460 tons	Trøndelag
Ilsoya	Frøya	3900 tons	Trøndelag
Gåsholmen	Frøya	2340 tons	Trøndelag
Storbrannøya	Frøya	1560 tons	Trøndelag
Lille Torsøy	Hitra	5200 tons	Trøndelag
Korsholman	Hitra	3120 tons	Trøndelag
Helsøya	Hitra	3900 tons	Trøndelag
Osholmen	Hitra	3120 tons	Trøndelag
Svellungen	Hitra	3120 tons	Trøndelag
Kåholmen	Hitra	4680 tons	Trøndelag
Heggvika	Hitra	2340 tons	Trøndelag
Grønnholmsundet	Hitra	1820 tons	Trøndelag
Sengsholmen	Hitra	1560 tons	Trøndelag
Veddersholmen	Bjugn	4680 tons	Trøndelag
Flatøya	Bjugn	2340 tons	Trøndelag
Breidvika	Osen	5460 tons	Trøndelag
Indre Skjervøy	Osen	7020 tons	Trøndelag
Drogsholmen	Roan	2340 tons	Trøndelag
Svefjorden	Osen	2340 tons	Trøndelag
Almurden	Flatanger	3900 tons	Trøndelag
Estenvika	Flatanger	2340 tons	Trøndelag
Austvika	Flatanger	3120 tons	Trøndelag
Bjørgan	Flatanger	5460 tons	Trøndelag
Dalavika	Flatanger	1560 tons	Trøndelag
Feøyvika	Flatanger	5460 tons	Trøndelag
Bragstadsundet III	Fosnes	3900 tons	Trøndelag
Kjelneset	Fosnes	4680 tons	Trøndelag
Ølhammaren	Fosnes	2340 tons	Trøndelag
Vedøya	Fosnes	3120 tons	Trøndelag

Table 6.2: The implemented growth model (Skretting, 2011)

Fish Class		Temperature (Celsius)									
f	V_f	0,5	1,0		14,0	14,5	15,0	15,5		19,5	20,0
1	0,03	0	0		0,04	0,04	0,04	0,04		0,03	0,03
2	0,05	0	0	0,05	0,05	0,05	0,05	0,05	0,04
3	0,07	0	0		0,07	0,07	0,07	0,07		0,06	0,06
....										
52	3,25	0,02	0,02		0,67	0,67	0,67	0,67		0,55	0,51
53	3,38	0,02	0,02		0,68	0,68	0,69	0,68		0,55	0,52
54	3,50	0,02	0,02		0,69	0,69	0,70	0,69		0,56	0,53
55	3,63	0,03	0,03		0,70	0,70	0,71	0,70		0,56	0,53
56	3,75	0,03	0,03		0,71	0,71	0,71	0,71		0,57	0,54
57	3,88	0,04	0,04		0,72	0,72	0,72	0,72		0,58	0,54
58	4,00	0,04	0,04		0,73	0,73	0,73	0,73		0,58	0,55
59	4,13	0,04	0,04		0,74	0,74	0,74	0,74		0,59	0,56
60	4,25	0,04	0,04		0,74	0,75	0,76	0,75		0,60	0,56
61	4,38	0,04	0,04		0,76	0,76	0,76	0,76		0,60	0,57
62	4,50	0,04	0,04		0,77	0,77	0,77	0,76		0,61	0,57
63	4,63	0,04	0,04	0,78	0,78	0,78	0,78	0,61	0,57
64	4,75	0,04	0,04		0,78	0,79	0,80	0,79		0,62	0,58
65	4,88	0,04	0,04		0,79	0,80	0,80	0,80		0,63	0,59
66	5,00	0,05	0,05		0,80	0,80	0,80	0,80		0,64	0,60
67	5,13	0,05	0,05		0,81	0,81	0,81	0,81		0,65	0,60
68	5,25	0,05	0,05		0,83	0,83	0,83	0,83		0,65	0,61
69	5,38	0,05	0,05		0,83	0,83	0,83	0,83		0,65	0,61
70	5,50	0,05	0,05		0,83	0,83	0,83	0,83		0,65	0,61
71	5,63	0,05	0,05		0,83	0,83	0,83	0,83		0,65	0,61
72	5,75	0,05	0,05		0,83	0,83	0,83	0,83		0,65	0,61
73	5,88	0,05	0,05		0,83	0,83	0,83	0,83		0,65	0,61
74	6,00	0,05	0,05		0,83	0,83	0,83	0,83		0,65	0,61
75	6,13	0,05	0,05		0,83	0,83	0,83	0,83		0,65	0,61
....										
80	6,75	0,05	0,05		0,83	0,83	0,83	0,83		0,65	0,61
81	6,88	0,05	0,05	0,83	0,83	0,83	0,83	0,65	0,61
82	8,00	0	0		0	0	0	0		0	0

Mid for the years 1998 to 2006. Each of the 9 years constitutes a growth scenario, as shown in figure 6.2. All scenarios are given an equal probability of $1/9$, as the amount of data is not sufficient to create a probability distribution. Since the resolution in the growth model is 0.5 degrees Celsius, all temperatures are rounded to the nearest half degree. Even though Region Mid is made up of two separate regions, Møre og Romsdal and Trøndelag, the same growth scenarios are used for all locations in the models. Ideally, since the locations are spread over several fjords and coastal areas and therefore may have variations in temperature, different temperature data should be used for certain locations, but individual location data was not available.




Figure 6.2: Growth scenarios, given by temperature data from Marine Harvest Region Mid for the years 1998 to 2006 (Marine Harvest, 2011)

6.2.5 Prices and Contracts

Since different sized fish achieve different prices in the market, prices must be included for each of the 7 sales classes j in the implemented models. The spot sales prices used are given by The Norwegian Seafood Federation, who records weekly spot prices in NOK for Atlantic salmon weighing 1-2, 2-3, 3-4, 4-5, 5-6, 6-7, and 7-8 kilograms. In the implemented models, positive prices are only included for the sales classes where fish weigh 3-4, 4-5, 5-6, and 6-7 kilograms, as these are normally the only classes sold from locations in Marine Harvest Region Mid. These weights are also the most common in the rest of Norway; in 2010, 86% of the salmon sold weighed between 3 and 7 kilograms. By setting the spot price for salmon weighing 1-2 and 2-3 kilograms equal to zero, fish belonging to these

sales classes will only be harvested if absolutely necessary in order to comply with biomass constraints or fallowing restrictions, referred to as emergency harvesting. Setting the spot price for salmon weighing 7-8 kilograms equal to zero ensures that the fish are harvested before they reach this size. The monthly prices used are calculated as the average of the weekly prices in each month.

The spot purchase price and the contract prices are both based on the given FHL spot price. The spot purchase price is modelled with a fixed 10 NOK markup from the spot price, based on the discussion in section 2.4 regarding why forced purchases in the spot market can be costly. Contract prices are not modeled with fixed markup, but are instead generated randomly to realistically create a set of 100 contracts with varying prices. The randomly created prices range between the spot price and 10 NOK above the spot price in the period of delivery. The premium included in the contract prices is needed in order to give the producer an incentive to enter contracts. At the same time, the premium cannot be larger than the spot purchase price mark-up, as this would introduce arbitrage. The topic of deciding spot purchase prices and contract prices should ideally be devoted more attention, but since price is assumed to be deterministic, this is not the focus of this thesis.

The models are tested using price data from 1997, 2001 and 2010, where each problem instance is chosen to represent a characteristic price development. In 1997, prices were steady all through the year, in 2001 prices decreased during the year, while in 2010 the prices were volatile and increasing. Note that since prices are assumed deterministic and the contract prices are generated randomly, studying the effects on the model results of using different price data is of limited value. Testing the models using three problem instances is therefore done mostly for illustrative purposes. The price development in problem instance 1, 2 and 3 is shown in figures 6.3, 6.4, and 6.5, respectively.

6.2.6 Additional Data

Temperatures, prices, and contracts are not the only elements in the data set. The initial biomass, the release plan, and the set of harvesting periods are synthetic, but realistic for Marine Harvest Region Mid. This is also the case for the set of fallowing periods, where two month fallowing is used. The set of previously entered contracts are created randomly in the same manner as for the set of contracts that can be entered. The cost of caring for the biomass is based on data from Skretting, while the processing ratios that account for the weight lost due to gutting during slaughtering are set to equal to 1, due to limited access to accurate information. Detailed data is included in the electronic documentation.




Figure 6.3: Problem instance 1: 1997 (Norwegian Seafood Federation, 2010)




Figure 6.4: Problem instance 2: 2001 (Norwegian Seafood Federation, 2010)

6.3 Two-stage Stochastic Model

The TS model is applied to a problem with two stages, where uncertainty is resolved after decisions regarding contract sales are made in the first period, Jan-




Figure 6.5: Problem instance 3: 2010 (Norwegian Seafood Federation, 2010)

uary. The scenario tree for the two-stage problem is shown in figure 6.6. The root node represents January, in which the information available is common for all of the scenario problems. The temperature in January is the average of all the January temperatures in the set of scenarios. Each node in period two, February, represents the first period in the second stage, where all information regarding the remaining periods is available. This implies that the non-anticipativity constraints only apply to decisions taken in the first period.

6.4 Multistage Stochastic Model

The MS model is applied to a problem where new information is available at two points in time, resulting in a total of three stages. First, uncertainty is resolved after the first period, January, but with only three possible outcomes, as opposed to the 9 outcomes after the first period in the two-stage problem. The second stage lasts for four periods, February to May, before uncertainty again is resolved after the fifth period. Again, there are three possible outcomes in each node, resulting in a total of 9 scenarios, the same as for the two-stage problem. Each node in period six, June, represents the first period in the third stage, where all information regarding the remaining seven periods is available. June is chosen as the start of the third stage since June to September is the period of the year where the variations in temperature are the greatest (see figure 2.3). Capturing this uncertainty has a value in the model. The scenario tree for the multistage problem is shown in figure 6.7.

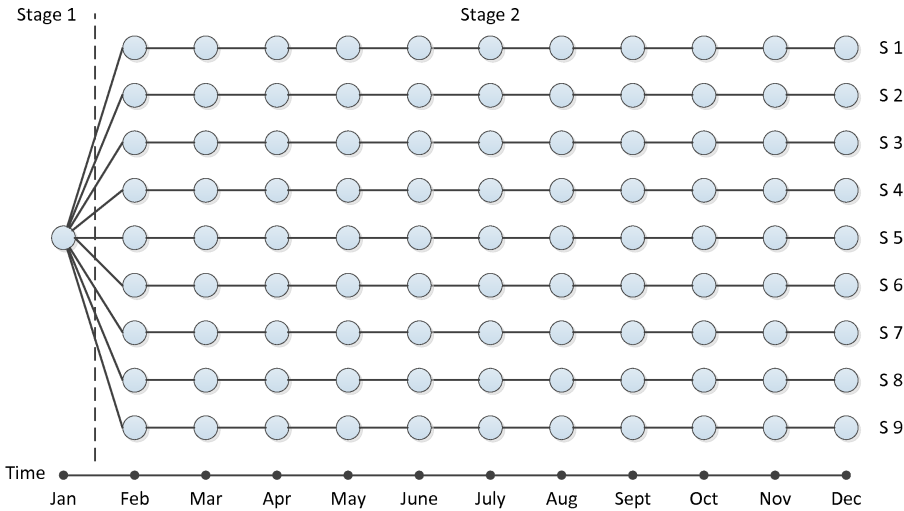


Figure 6.6: The scenario tree for the two-stage problem

Again, the temperature in first stage, January, is the average of all the January temperatures in the set of scenarios. For the second and third stages, the information structure is not as intuitive. In order to have three outcomes in the second stage, three groups of three temperature scenarios must be chosen, giving a total of 280 possible combinations¹. An important aspect of deciding which combination to use is that there is no intertemporal temperature correlation, as discussed in section 2.4, and only one stochastic variable. This means that there are no statistical properties in the problem that must be taken into consideration when deciding on the combination to be used. Therefore, one way of grouping the scenarios is by randomly creating three groups. The problem with this approach is that for each group, the temperature in each period in the second stage is the average temperature of the three scenarios in the group. Taking the average removes some of the variation in data and results in three outcomes that are likely to be very similar, thus limiting the point of resolving uncertainty. Therefore, to emphasize the uncertainty in biomass development, the scenarios are grouped in three groups of three so that the total standard deviation (standard deviation summed over the three groups) is minimized, thereby also minimizing the effect of using average values. Using this approach results in a low growth outcome, a medium growth outcome, and a high growth outcome.

In the TS model, the non-anticipativity constraints only apply to contract sales, since these are the only decisions taken in the first period. In the MS model, decisions made regarding slaughtering, spot sales and spot purchases in all periods in the second stage must also be included in the non-anticipativity constraints. In this way, the decisions taken in the three scenario problems in a group in

¹ $9!/(3!*3!*3!*3!) = 280$ possible combinations

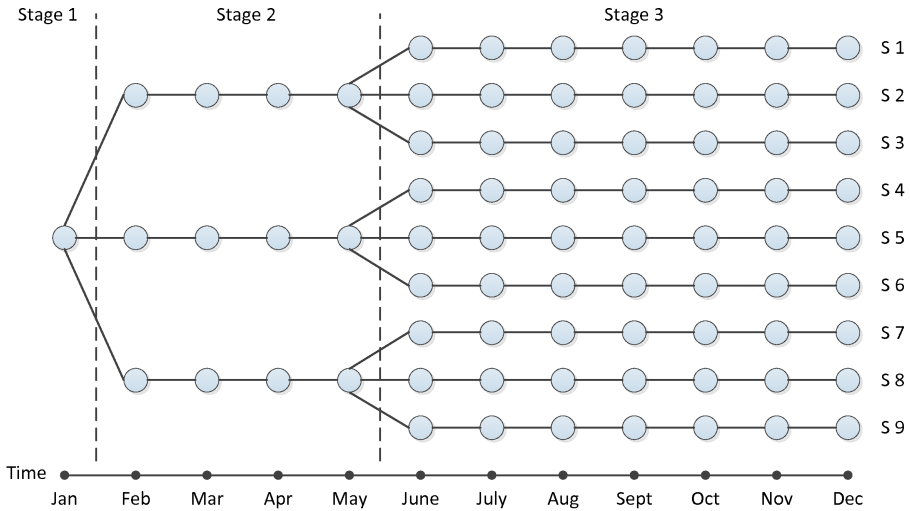


Figure 6.7: The scenario tree for the multistage problem

the second stage are guaranteed to be equal. The information structure ensured by the non-anticipativity constraints in the first five periods is clear in figure 6.8.

The scenario tree in 6.7 is the result of decisions taken regarding how to represent the information and decision structure in the problem. Ideally, the number of stages and scenarios in a multistage model is determined based on the problem characteristics. However, as the problem size grows exponentially, realistically portraying the information process is often difficult. For the implemented models, a limit of approximately 10-15 scenarios are manageable on a personal computer, depending on the data set. This limits the number of ways the scenario tree can be structured, and necessitates certain tradeoffs. The advantage of the current scenario tree is that having three outcomes in the second stage allows for an intuitive low/medium/high structure. At the same time, a total of 9 scenarios allows for utilizing all of the 9 years of temperature data available.




Figure 6.8: Scenarios in the multistage problem. Note that the temperature in January is equal for all scenario problems, and the temperature follows one of three paths from February to May.

Chapter 7

Results and Discussion

This chapter presents the results of applying the three models introduced in the previous chapter to Marine Harvest Region Mid. Section 7.1 summarizes the main results and explains how the results are calculated, while more detailed results for the multistage problem are presented in section 7.2. The results are then discussed in section 7.3, along with the consequences of varying key input parameters.

7.1 Main Results

The solution values obtained from the deterministic (DET), the two-stage stochastic (TS), and the multistage stochastic (MS) models can be described using the concepts introduced in chapter 3. Before presenting the main results, some of these concepts and relevant notation will be repeated as part of the explanation of how the results are calculated.

7.1.1 EV, RP, WS and EEV

The main results can be presented as the set of solution values obtained when solving the problem instances with the three different models. The solution given by the DET model is the expected value solution, EV, calculated by replacing the set of scenario problems with a single expected value scenario problem. The solutions given by the TS and MS models are the here-and-now solutions, denoted by RP. RP is obtained by solving the set of scenario problems, with non-anticipativity constraints ensuring that decisions in each scenario problem are taken with respect to the information structure. Then, if the non-anticipativity constraints in the TS/MS models are relaxed, the solution obtained is the wait-and-see solution, denoted by WS.

The expected value of using the EV solution when uncertainty is included is denoted by EEV. As opposed to EV and RP, which are given directly from their respective models, EEV requires some calculation. For the two-stage model, EEV is found by solving the TS model with the first stage decisions fixed to the decisions taken in the DET model. Since the first stage consists of only the first period, where decisions regarding which contracts to enter are the only ones taken, calculating EEV for the two-stage problem is straight forward.

In problems with more than two stages, calculating EEV is a more complicated matter, and requires applying one of the two methods proposed in 3.3. In the first and simplest method, approach A, EEV is found by using the MS model with some or all of the decision variables in the first two stages fixed at the optimal values obtained in the DET model solution, EV. Approach A is an extension of the method used for two-stage problems, modified so that the non-anticipativity constraints are taken into account. In the second method, approach B, the scenario tree is divided into sub-trees, and a set of expected value problems are solved. This method has the advantage that it portrays more accurately the information structure present in the problem, thereby giving a more realistic EEV value. Unfortunately, only approach A can be applied to the current model, since approach B is sensitive to the errors introduced by having to round average temperatures in the set of expected value problems. The source of the problem is the implemented growth model which has a resolution of 0.5 degrees Celsius, meaning that all average temperatures must be rounded to the nearest half degree. In approach A, this is not a problem, since only one expected value problem is solved. In addition, the rounded values used in approach A are averages taken over all scenarios, resulting in an error that is insignificant. With approach B, however, a total of four expected value problems must be solved, and three of the problems involve using the rounded average values of only three scenarios. The resulting errors are significant enough that EEV cannot be calculated using approach B.

Which decision variables to fix when calculating EEV using approach A depends on the problem characteristics, and requires a careful analysis of the set of decisions. As discussed in 3.2.1, delayed decisions which are meant to compensate for, exploit, or possibly correct negative effects resulting from decisions taken earlier are referred to as recourse decisions. Understanding which decisions are recourse decisions is important because these decisions are normally not fixed when calculating EEV using approach A. In the current models, two of the four decisions are examples of typical recourse decisions: spot sales and spot purchases. Spot sales and purchases are carried out based on previous decisions regarding contract sales, they are taken one period at a time only for the current period, and they do not directly affect decisions that are to be taken in later periods. On the contrary, deciding which contracts to enter is a typical example of a decision which is not a recourse decision. For the fourth decision, slaughtering, the recourse characteristics are not as apparent as for the first three. On the one hand, slaughtering is a result of previous decisions regarding contract sales, and are taken one period at a time for only the current period. On the other hand,

slaughtering affects the biomass in the next period, thereby influencing subsequent decisions. Regardless of whether slaughtering should be classified as a recourse decision or not, slaughtering decisions cannot be fixed in the MS model, since fixing slaughtering decisions leads to infeasible solutions. The reason for this is that slaughtering decision that can be taken in the single expected scenario problem are not necessary feasible in the MS model, since the fish that are to be slaughtered may not exist due to uncertain biomass development.

7.1.2 Summarized Results

The values EV, RP, WS and EEV are now given for the three problem instances presented in the previous chapter, for both the two-stage and multistage problems. Price data from 1997, 2001, and 2010 is used in problem instance one, two, and three, respectively, while all other data in the data set is equal. As explained in chapter 4, most costs are omitted, meaning that the results are only of comparative value and must not be interpreted as the company's net profit. Tables 7.1 and 7.2 show the results for the two-stage and multistage problems, respectively.

Table 7.1: Results for the two-stage problem (1000 NOK)

Year	EV	EEV	WS	RP
1997	1 540 930	1 537 200	1 538 080	1 537 800
2001	1 397 840	1 393 600	1 393 730	1 393 620
2010	2 456 910	2 438 430	2 439 920	2 438 620

Table 7.2: Results for the multistage problem (1000 NOK)

Year	EV	EEV	WS	RP
1997	1 543 270	1 536 530	1 548 220	1 546 170
2001	1 397 990	1 398 390	1 402 240	1 400 810
2010	2 451 300	2 442 630	2 454 380	2 452 040

In all three problem instances for both the two-stage and multistage models, the WS solution is greater than the RP solution, which in turn is greater than the EEV solution. This coheres with property 3.3.3, which is valid for all recourse problems. Also note that the EV solution is not the same for the two-stage and multistage problems. This is due to the errors introduced by having to round average temperatures in DET model, as discussed above.

7.2 Detailed Results

A more in-depth presentation of the results is now given for the multistage problem for each of the three problem instances. First, end of period biomass, spot sales, spot purchases, and contract sales are given per period for the RP and EEV solutions. The RP and EEV solutions are presented to provide a comparison of the deterministic and stochastic solutions. For the two-stage problem, the differences in the RP and EEV results are not significant enough to be displayed graphically, and the detailed results are instead included numerically in the electronic documentation. Also given in the electronic documentation are the complete results for the EV, RP, WS and EEV solutions, for all three problem instances, for both the two-stage and multistage problems. The complete results specify the amount (in number of fish and in kilograms) of fish harvested, sold, purchased, and released, per fish or sales class, per location, per period.

Figures 7.1, 7.2, and 7.3 show the detailed results for problem instances 1, 2, and 3, respectively. In each figure, the left column displays the details of the RP solution, while the right column displays the details of the EEV solution. Both biomass and sales are given in tons per period. Each colored line in the charts is a single scenario. In the end of period biomass charts, the combined MAB for the two regions Møre og Romsdal and Trøndelag is given by the dashed line. The same results are given per region in the electronic documentation.

Several aspects of the results in the figures are worth commenting. Firstly, common for all three problem instances and for both the RP and EEV solutions is that the end of period biomass decreases in May and December. The reason for the biomass decrease starting in May is that low winter temperatures combined with continues slaughtering throughout the year makes it difficult to the fully exploit the MAB limits in late winter and spring. The decrease in December is caused by the way that end-of-horizon conditions are dealt with in the models, and would not be present with an increased planning horizon.

Another important observation is that for problem instances 1, the amount of spot purchases in the EEV solution is considerably larger than for the RP solution (where minimal spot purchases only occur in one scenario). This illustrates how decisions in a stochastic model can be made so that costly “penalties” are avoided. Problem instance 2 illustrates the opposite. Here, the amount of spot purchases is higher in the RP solution than in the EEV solution. This occurs when profits achieved in certain scenarios from entering a contract are sufficiently large that even with spot purchasing occurring in other scenarios, entering the contract is still more profitable than spot sales. Problem instance 3 illustrates that uncertainty is not always relevant for spot purchases, since for this problem instance the amounts of spot purchases in the EEV and RP solutions are equal. Regarding spot sales, no conclusions should be drawn since the amount sold spot in each period is a result of randomly generated contracts. Also note that in the charts showing the EEV solution contract sales, only a single scenario is present from January to May due to the use of approach A, as discussed above.



Figure 7.1: Detailed results for problem instance 1. The left column displays the details of the RP solution, the right column displays the details of the EEV solution. For all charts, periods are given on the x-axis, while the y-axis shows biomass in tons. The colored lines represent individual scenarios.



Figure 7.2: Detailed results for problem instance 2. The left column displays the details of the RP solution, the right column displays the details of the EEV solution. For all charts, periods are given on the x-axis, while the y-axis shows biomass in tons. The colored lines represent individual scenarios.

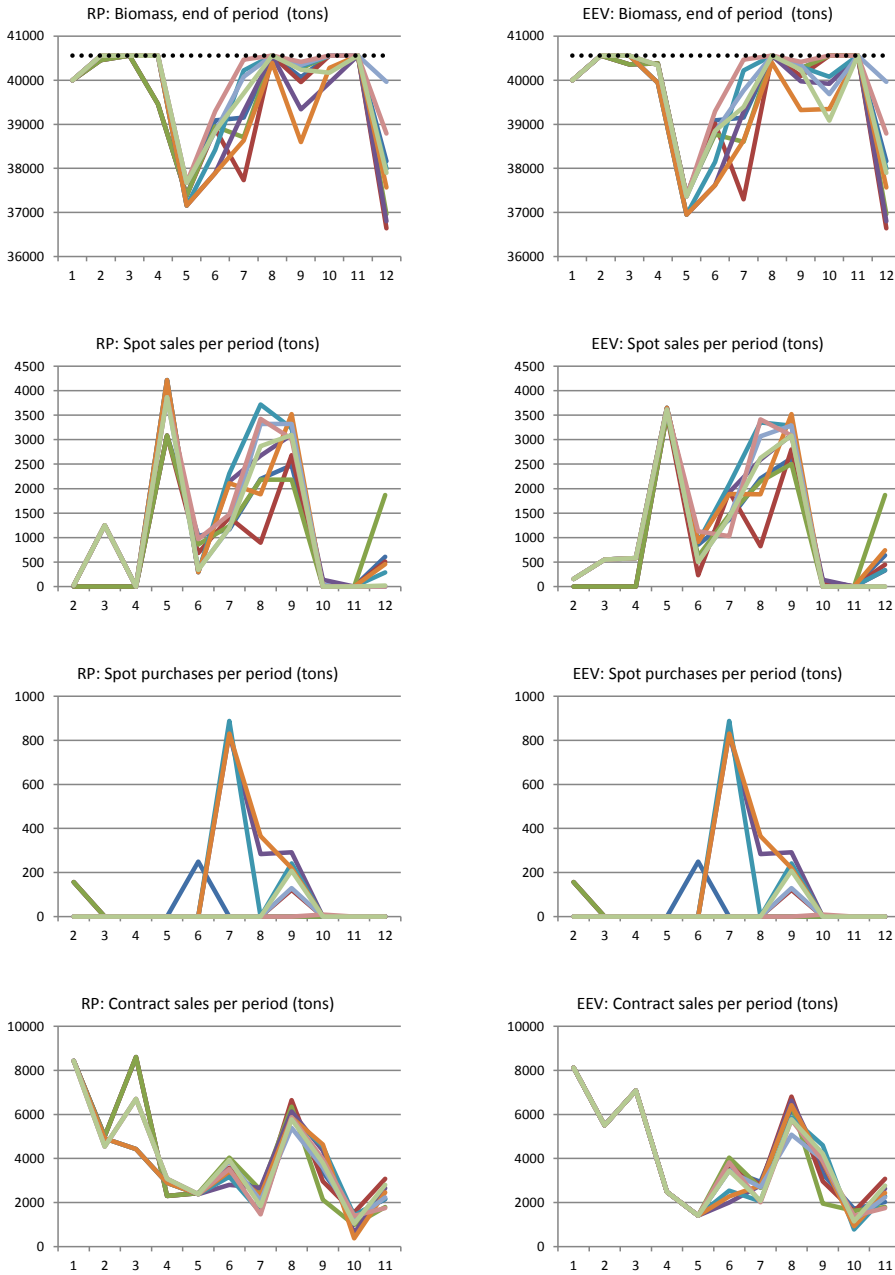


Figure 7.3: Detailed results for problem instance 3. The left column displays the details of the RP solution, the right column displays the details of the EEV solution. For all charts, periods are given on the x-axis, while the y-axis shows biomass in tons. The colored lines represent individual scenarios.

In stochastic models with rolling horizons, first stage decisions are of special interest. They are the only decisions which are implemented after running the model, while implementation of decisions belonging to later stages is postponed until new information is available. In both the two-stage and multistage problems, the only decisions made in the first stage are regarding which contracts to enter. The first stage decisions taken in the multistage problem are presented here, while results for the two-stage model are included in the electronic documentation.

Tables 7.3, 7.4, and 7.5 give the first stage decisions in the RP and EV/EEV solutions for problem instances 1, 2, and 3, respectively. Note that the first stage decisions in the EV and EEV solution are the same, since EEV is calculated by fixing the first stage decisions to those given by EV solution. The leftmost column in the tables, *Cont.*, gives the contract indexes, while the next column, *Price*, gives the price of each contract in NOK per kilogram. The columns *Class* and *Del.* give the sales class and the delivery period for each contract, respectively. The maximum quantities for each contract are given in the column *Max*. Finally, the two rightmost columns give the tons of salmon sold in each contract in the RP and EV/EEV solutions.

Though the results for the RP and EV/EEV solutions are very similar for all three problem instances, the solutions are not equal. This demonstrates that the decisions taken by the DET model are not always optimal when uncertainty is taken into account.

Table 7.3: First stage decisions in the RP and EV/EEV solutions for problem instance 1

Cont.	Price	Class	Del.	Max	RP	EV
14	37	5	8	230	230	230
19	26	3	2	810	0	0
30	33	6	5	1100	1100	1100
41	27	4	9	470	0	0
47	29	3	6	330	0	0
48	27	5	10	360	0	360
60	32	3	5	800	69	50
89	29	6	6	1000	1000	1000
100	38	6	9	920	920	920

Table 7.4: First stage decisions in the RP and EV/EEV solutions for problem instance 2

Cont.	Price	Class	Del.	Max	RP	EV
11	32	5	5	1500	1500	1500
30	32	6	5	1100	1100	1100
64	25	5	11	1340	1340	1340
69	18	4	12	1050	0	0
72	33	3	5	1370	92	51
86	27	3	3	1000	0	0
92	31	4	2	780	780	780
97	29	3	8	200	0	0
100	31	5	5	260	260	260

Table 7.5: First stage decisions in the RP and EV/EEV solutions for problem instance 3

Cont.	Price	Class	Del.	Max	RP	EV
10	48	6	4	280	280	280
17	40	5	3	980	980	980
18	44	5	5	1240	0	313
19	49	5	4	620	620	620
30	49	6	5	1100	1100	1100
39	52	5	12	590	590	590
42	49	6	7	300	300	300
46	36	3	10	380	0	0
49	42	5	10	1460	0	0
54	44	4	2	1350	1350	1220
84	45	3	9	610	610	122
87	43	6	2	1380	1380	1380
88	51	5	12	1220	1220	1220
89	49	3	12	1240	0	0

7.3 Evaluation and Analysis

7.3.1 VSS and EVPI

Continuing to use the concepts introduced in chapter 3, the TS and MS models can be evaluated based on the main results presented in the previous section. Two methods of evaluation have been presented; the expected value of perfect information, EVPI, and the value of the stochastic solution, VSS. EVPI compares the WS and RP solutions, measuring the amount a decision maker would be ready to pay in return for complete information about the future. Attaining complete information regarding future seawater temperatures is not possible, making EVPI a purely theoretical measure for the current problems. VSS compares the RP and EEV solutions, giving the expected value of using a stochastic model over a deterministic model. Tables 7.6 and 7.7 give VSS and EVPI for each problem instance for the two-stage and multistage problems, respectively.

Table 7.6: Evaluation of the TS model (1000 NOK)

Year	EVPI	EVPI (%)	VSS	VSS (%)
1997	280	0,018 %	600	0,039 %
2001	110	0,008 %	20	0,001 %
2010	1 300	0,053 %	190	0,008 %

Table 7.7: Evaluation of the MS model (1000 NOK)

Year	EVPI	EVPI (%)	VSS	VSS (%)
1997	2 050	0,133 %	9 640	0,623 %
2001	1 430	0,102 %	2 420	0,173 %
2010	2 340	0,095 %	9 410	0,384 %

For both the TS and MS models, EVPI and VSS are positive for all problem instances in agreement with property 3.3.3. Also, for all problem instances, both EVPI and VSS are lower for the TS model than for the MS model. This is as expected, since less uncertainty is included in the two-stage problem than in the multistage problem. The almost zero VSS and EVPI for the TS model is largely explained by the fact that the contracts sold before uncertainty is resolved amounts to only 1/12 of the total contract decisions taken, thereby having limited impact. For the remaining 11 periods, the problem is solved deterministically, explaining the limited difference between the RP and EEV solutions and RP and WS solutions for the two-stage problem. The results clearly show that using a stochastic model with only two stages provides almost no additional value compared to the much simpler deterministic approach.

Though still quite low, VSS and EVPI are much higher for the multistage problem than for the two-stage problem. Problem instance 1 has the highest VSS, where using the stochastic solution increases profits by almost 10 million NOK. The increase in profit must be considered together with the additional computational effort needed when choosing whether to use a stochastic or deterministic approach.

7.3.2 Spot Purchase Price and Contract Price Markups

In the results presented so far, the spot purchase price is fixed at 10 NOK above the FHL spot price, and contract prices are generated randomly, ranging from the spot price to 10 NOK above the spot price in the period of delivery. The influence of contract and spot purchase markups on the optimal solution are now further investigated by varying the markups in the multistage problem while keeping the remainder of the data set fixed.

Spot Purchase Price

The results of varying the spot purchase price markup between 0 and 20 NOK are shown in table 7.8, where the FHL spot price used is the same as in problem instance 3 (2010 prices). The solutions are evaluated in terms of VSS and EVPI in table 7.9.

Table 7.8: Results of varying the spot purchase price markup (1000 NOK)

Markup	EV	EEV	WS	RP
0	2 627 630	2 638 080	2 638 110	2 638 080
5	2 485 970	2 488 440	2 492 950	2 492 070
10	2 451 300	2 442 630	2 454 380	2 452 040
15	2 451 300	2 440 310	2 454 380	2 451 440
20	2 451 300	2 438 520	2 454 380	2 451 300

Table 7.9: Spot purchase price markup evaluation (1000 NOK)

Markup	EVPI	EVPI (%)	VSS	VSS (%)
0	30	0,001 %	-	0,000 %
5	880	0,035 %	3 630	0,146 %
10	2 340	0,095 %	9 410	0,384 %
15	2 940	0,120 %	11 130	0,454 %
20	3 080	0,126 %	12 780	0,521 %

The results in table 7.9 clearly illustrate that as the spot purchase price markup increases, so does both the value of using a stochastic model and the expected

value of perfect information. This is natural since the purchase price markup can be thought of as the “cost” of compensating or correcting for uncertainty. As this cost increases, it becomes more important to take the uncertainty into consideration. Note that when the spot purchase price markup is zero (the spot purchase price equals the spot sales price), there is no penalty for having to buy fish in the spot market in order to fulfill contract commitments. This results in both an EVPI and VSS of zero (the EVPI value of 30 is due to the mentioned rounding errors.) Also note that with a markup of 10, 15 or 20, the EV and WS solutions do not change since these solutions avoid all spot purchases and are therefore unaffected. The EV and WS solutions change when the markup is 0 and 5 because of the arbitrage introduced when contracts priced above the spot purchase price make spot purchases profitable.

Contract Price

The effects of the contract price markup on the optimal solution can be studied by using a set of fixed contracts (minimum and maximum amount, delivery period, sales class, and entering period are kept unchanged) with different markup values. The results of varying the contract price markup from 0 to 15 in the multistage problem is given in table 7.10. The solutions are evaluated in terms of VSS and EVPI in table 7.11.

Table 7.10: Results of varying the contract price markup (1000 NOK)

Markup	EV	EEV	WS	RP
0	2 157 180	2 165 250	2 167 660	2 167 630
5	2 299 850	2 297 210	2 307 140	2 305 300
7,5	2 375 380	2 369 420	2 380 500	2 378 060
10	2 451 300	2 442 630	2 454 380	2 452 040
15	2 624 910	2 617 540	2 626 590	2 623 860

Table 7.11: Contract price markup evaluation (1000 NOK)

Markup	EVPI	EVPI (%)	VSS	VSS (%)
0	30	0,001 %	2 380	0,110 %
5	1 840	0,080 %	8 090	0,351 %
7,5	2 440	0,103 %	8 640	0,363 %
10	2 340	0,095 %	9 410	0,384 %
15	2 730	0,104 %	6 320	0,241 %

Table 7.11 shows a similar pattern for contract price markup as for spot purchase price markup, except for the inconsistency in EVPI when the markup is 7,5, and the inconsistency in both the EVPI and VSS when the markup is 15. A higher

EVPI with a markup of 7,5 than with a markup of 10 cannot be explained by rounding errors. It can therefore not be concluded based on the EVPI results above that EVPI increases with increasing contract price mark-up. In regards to the inconsistency when the markup is 15, this can be explained by the introduction of arbitrage when the spot purchase price is lower than some of the contract prices. Clearly, as long as arbitrage is absent, VSS increases with increasing contract price markup. Note that when the contract price markup is zero (contract prices are equal to the spot price), all sales should be made in the spot market so that no flexibility is lost. The reason that VSS is not zero in this situation is that there is nothing in the EV model that makes spot sales preferable over contract sales if the two prices are equal. If this information were to be added, both the EVPI and VSS would be zero (again, the EVPI value of 30 is due to the mentioned rounding errors.)

The complete results for the EV, RP, WS and EEV solutions for all of the markup problem instances are included in the electronic documentation. The results specify the amount (in number of fish and in kilograms) of fish harvested, sold, purchased, and released, per fish or sales class, per location, per period.

Chapter 8

Concluding Remarks

Understanding the uncertainty present in salmon farming is an important part of successfully planning harvesting and future sales. This thesis describes the most important sources of uncertainty after the salmon are released in seawater, and presents stochastic programming as a tool for including this uncertainty in a problem formulation. The resulting multistage stochastic model provides salmon producers with a tool that can aid them in making profitable decisions regarding harvesting and future sales. The model considers both the uncertainty in biomass development and future salmon prices, as well as the many constraints that together define salmon aquaculture.

Applying a set of simplified models to Marine Harvest Region Mid illustrates how the model can be implemented and used, as well as how a quantitative assessment of the gains from implementing a stochastic solution can be performed. The evaluation of the stochastic models indicates that with only two stages, the stochastic solution is almost equal to the deterministic solution. For the multi-stage stochastic model, the value of the stochastic solution is substantially higher, though still marginal, largely due to the simplifications made in the implementation. The results also show that as the spot purchase price markup or contract price markup increases, so does the value of using a stochastic model.

The implemented models are subject to several simplifications which are necessary in order to make the models computationally tractable. Further treatment of these simplifications provides a natural extension of the work done in this thesis. The two main simplifications are that price is assumed deterministic and that the actual information and decision process is reduced to only being modelled by two or three stages.

Even though Fish Pool provides salmon farmers with financial tools that can be used to partially reduce price risk, this does not prevent the model from being significantly weakened by omitting price uncertainty. Including price uncertainty is recommended if the model is to be applied to a real planning problem. Adding

price stochasticity requires a discussion of the resolution that is to be used in the model so that the short term price volatility is considered to a sufficient degree. A possible, but rather comprehensive extension to the model, would be to implement the financial salmon contracts offered by Fish Pool as methods for hedging price volatility. This extension requires a detailed discussion of financial tools and risk management, and is thus outside the scope of this thesis.

A second possible extension is to increase the number of stages in the stochastic implementations in order to more accurately portray the actual decision and information process inherent when planning sales and harvesting. Studying the effects on problem size, computational effort, and the model results would provide a more satisfying basis for evaluating whether using a stochastic model is appropriate.

Including price uncertainty and increasing the number of stages will introduce additional decision variables and scenarios, thereby increasing the size and complexity of the implemented models. The suggested extensions will therefore likely require aggregation of periods, locations, or fish classes, or the use of more advanced solution methods such as decomposition. Access to additional computer power would allow for a more thorough study of the effects of incorporating uncertainty in a stochastic model.

In addition to the suggested extensions, the model can benefit from adopting a more complete value chain perspective. A completely integrated model including all activities from breeding to final processing would allow for better coordination of planning between different departments. Incorporating growing and release of smolt in the model is especially relevant since many companies are vertically integrated, owning and controlling both the freshwater and seawater parts of the value chain. Expanding the model to include the freshwater activities would affect sales planning by improving the control over the availability of smolt of different sizes, thereby increasing the flexibility in the timing of smolt releases.

Bibliography

- AKVA Group (2010). Cage farming aquaculture.
- Alnæs, A. and Skagen, M. (2009). Risikostyring i laksemarkedet - en studie av fish pools derivater. Master's thesis, NHH.
- Andersen, T. and Tveterås, S. (2008). The price responsiveness of salmon supply in the short and long run. *Marine Resource Economics*, 23(4):425–37.
- Asche, F. (2010). Conferance call with Frank Asche, University of Stavanger.
- Asche, F. and Bjørndal, T. (2005). Demand structure for fish.
- Asche, F. and Guttormsen, A. (2002). *Patterns in the relative price for different sizes of farmed fish*. Institute for Research in Economics and Business Administration.
- Asche, F. and Salvanes, K. (1997). Market delineation and demand structure. *American Journal of Agricultural Economics*, 79(1):139–150.
- Asche, F. and Salvanes, K. (1998). The demand for salmon in the European Union: the importance of product form and origin. *Canadian Journal of Agricultural Economics/Revue canadienne d'agroeconomie*, 46(1):69–81.
- Asche, F. and Tveterås, R. (2007). Testing cost vs. profit function. *Applied Economics Letters*, 14(10):715–718.
- Baker, K. (1977). An experimental study of the effectiveness of rolling schedules in production planning. *Decision Sciences*, 8(1):19–27.
- Barros, C. (2011). Presentation: recovering from ISA: the story of the Chilean salmon industria.
- Bergfjord, O. (2009). Risk perception and risk management in norwegian aquaculture. *Journal of Risk Research*, 12(1):91–104.
- Birge, J. and Louveaux, F. (1997). *Introduction to stochastic programming*. Springer Verlag.

- Cermaq ASA (2011). Cermaq fish farming. <http://www.cermaq.com/portal/wps/wcm/connect/cermaqen/home/fishfarming/>. Year denotes when the data was accessed.
- Directorate of Fisheries (2010). Fiskeridirektoratets statistikkbank. Contact: Merete, F.
- Escudero, L., Garín, A., Merino, M., and Pérez, G. (2007). The value of the stochastic solution in multistage problems. *Top*, 15(1):48–64.
- Fauske, M. (2011). Telephone conversation with Merete Fauske, Directorate of Fisheries.
- FHL Havbruk (2005). Akvakultur i Norge. Technical report, Norwegian Seafood Federation.
- FICO (2010). *Getting Started with Xpress*. Fair Isaac Corporation. Release 7.1.
- Guttormsen, A. and Forsberg, O. (2004). The value of information in salmon farming: harvesting the right fish to the right time. *Discussion Papers*.
- Higle, J. (2005). Stochastic programming: optimization when uncertainty matters. *Tutorials in Operations Research*, pages 1–24.
- Institute of Marine Research (2009). Kyst og havbruk 2009.
- Just, R. and Pope, R. (1978). Stochastic specification of production functions and economic implications* 1. *Journal of Econometrics*, 7(1):67–86.
- Kall, P. and Wallace, S. (1994). *Stochastic programming*. John Wiley & Sons Inc.
- Kontali Analyse (2010). Salmon world 2010.
- Kumbhakar, S. and Tveterås, R. (2003). Risk preferences, production risk and firm heterogeneity*. *Scandinavian Journal of Economics*, 105(2):275–293.
- Lorentzen, T. (2010). Økt sjøtemperatur påvirker lakseoppdretten. *Klima: Norsk magasin for klimaforskning*, (2):31–33.
- Marine Harvest (2010a). Marine Harvest capital markets day.
- Marine Harvest (2010b). Marine Harvest Norway. <http://www.marineharvest.com/en/About-Marine-Harvest/Business-units/Marine-Harvest-Norway>. Year denotes when the data was accessed.
- Marine Harvest (2010c). Salmon farming industry handbook 2010.
- Marine Harvest (2011). Data received from Marine Harvest. Contact: Osnes, E.
- Midthun, K. (2007). *Optimization models for liberalized natural gas markets*. Norwegian University of Science and Technology, Faculty of Social Science and Technology Management, Department of Sociology and Political Science.

- Midthun, K. (2010). Lecture notes: introduction to stochastic programming.
- Ministry of Fisheries and Coastal Affairs (2004). Laksetildelingsforskriften. http://www.lovdata.no/for/sf/fi/xi-_20041222-_1798.html.
- Ministry of Fisheries and Coastal Affairs (2008). Akvakulturdriftsforskriften. http://www.lovdata.no/for/sf/fi/xi-_20080617-_0822.html.
- Ministry of Fisheries and Coastal Affairs (2009). Forskrift om bekjempelse av lus i akvakulturanlegg. http://www.lovdata.no/for/sf/fi/xi-_20090818-_1095.html.
- Norwegian Seafood Federation (2010). FHL statistics. Contact: Grøttum, J.A.
- Oglend, A. and Sikveland, M. (2008). The behaviour of salmon price volatility. *Marine Resource Economics*, 23(4):487–506.
- Philpott, A. (2011). Stochastic programming introduction. <http://stoprog.org/index.html?spintroduction.html>. Year denotes when the data was accessed.
- Prékopa, A. (1995). *Stochastic programming*. Springer.
- Rockafellar, R. and Wets, R. (1991). Scenarios and policy aggregation in optimization under uncertainty. *Mathematics of operations research*, pages 119–147.
- Sandberg, F. (2010). Key figures from the aquaculture industry 2009. Technical report, Directorate of Fisheries.
- Sandmo, A. (1971). On the theory of the competitive firm under price uncertainty. *The American Economic Review*, 61(1):65–73.
- Sen, S. and Higle, J. (1999). An introductory tutorial on stochastic linear programming models. *Interfaces*, pages 33–61.
- Sethi, S. and Sorger, G. (1991). A theory of rolling horizon decision making. *Annals of Operations Research*, 29(1):387–415.
- Skretting (2011). Skrettings fôrkatalog 2011.
- Østen, J. (2010). Hvordan rømmer fisken.
- Sætre, R., editor (2007). *The Norwegian coastal current - oceanography and climate*. Tapir Academic Press.
- Sætre, R. (2011). Telephone conversation with Roald Sætre, Institute of Marine Research.
- The Research Council of Norway (2005). Forskningsbehov innen dyrevelferd i norge. Technical report, The Research Council of Norway.

- Trollvik, J. (2010). Global produksjon av laks. Technical report, Norwegian Seafood Export Council.
- Tveterås, R. (1999). Production risk and productivity growth: some findings for Norwegian salmon aquaculture. *Journal of Productivity Analysis*, 12(2):161–179.
- van der Vlerk, M. (2011). Stochastic programming lecture notes. <http://stoprog.org/>. Year denotes when the data was accessed.
- Wallace, S. (2003). Decision making under uncertainty: the art of modelling. *Molde University College*, 15(03).