Master's thesis

NTNU
Norwegian University of Science and Technology
Faculty of Economics and Management
Dept. of Industrial Economics and Technology Management

Randulf Martin Aasen

Tactical Production Planing for Atlantic Salmon Farming Under Uncertainty of Salmon Lice

Master's thesis in Industrial Economics and Technology Management
Supervisor: Peter Schütz
June 2021

NTNU
Kunnskap for en bedre verden

Randulf Martin Aasen

Tactical Production Planing for Atlantic Salmon Farming Under Uncertainty of Salmon Lice

Master's thesis in Industrial Economics and Technology Management
Supervisor: Peter Schütz
June 2021

Norwegian University of Science and Technology
Faculty of Economics and Management
Dept. of Industrial Economics and Technology Management

Preface

This thesis is written in connection with the subject TIØ4905 – Managerial Economics and Operations Research, Master's Thesis at the Norwegian University of Science and Technology during the spring semester of 2021. I would like to thank the industrial partners Aquagen AS, Eidsfjord Sjøfarm AS, and Sisomar AS for contributing and giving valuable insights to this thesis. I would also like to thank my academic advisor Peter Schütz for his invaluable guidance, especially his willingness to read and thoroughly critique drafts.

Abstract

In this thesis we develop an optimization model for tactical production planning during the seawater stage of the Atlantic salmon farming value chain. We develop a stochastic model that accounts for uncertainty in salmon lice and its impact on operations. We frame our model as a multi-stage stochastic mixed integer linear program. Due to computational constraints, this program is simplified into a two-stage stochastic program. The first-stage solution of this program is evaluated through a rolling horizon heuristic based simulation framework. We use this framework to investigate the potential benefits of introducing novel smolt types. Additionally, we assess the value of accounting for uncertainty when performing tactical production planning.

We optimize for two different objectives, namely raw harvest volumes, and harvest volumes weighted by relative prices. Our results show that introducing novel smolt types when maximizing for raw harvest volumes can increase the objective values by approximately 9.4 %, while we see an increase of 10.3 % when maximizing with respect to relative prices. We see a small shift towards harvesting heavier weight classes during late summer when optimizing for relative prices instead of raw harvest volumes. Moreover, we find that accounting for uncertainty during production planning has few benefits when it comes to increasing objective values, however, incorporating uncertainty tends to produce more robust production plans under some circumstances.

Sammendrag

I denne masteroppgava utvikler vi en optimeringsmodell for taktisk produksjonsplanlegging under saltvannsdelen av lakseoppdrettverdikjeden. Vi utvikler en stokastisk modell som tar hensyn til usikkerheten i lakselusnivåer og lusas operasjonelle betydning. Vi formulerer modellen som et flerstegs stokastisk lineært blandet heltallsprogram. Grunnet begrenset regnekraft forenkler vi dette programmet til et tostegs stokastisk program. Vi evaluerer førstestegsløsningen til dette programmet gjennom et rullerende horisont heuristikk basert simuleringsrammeverk. Vi bruker så dette rammeverket til å undersøke fordelene av å introdusere nye smolt typer. I tillegg undersøker vi verdien av å hensynta usikkerhet når man lager taktiske produksjonsplaner.

Vi optimerer med hensyn til to forskjellige målfunksjoner. Nemlig, rå slaktevolumer og slaktevolumer vektet med relativ pris etter vektklasse. Resultatene våre viser at å introdusere nye smolt typer kan øke målfunksjonen med omtrent 9,4 % når vi maksimerer med hensyn på rå slaktevolumer, og med 10,3 % når vi maksimerer med hensyn til relative priser. Vi ser også et lite skifte mot å slakte tyngre vektklasser på sensommeren når vi optimerer for relative priser istedenfor rå slaktevolumer. Dog finner vi at å hensynta usikkerhet når man lager produksjonsplaner har liten verdi når det kommer til å øke målfunksjonsverdiene, men at det sørger for at man i enkelte tilfeller får mer robuste produksjonsplaner.

Contents

List of Figures

List of Tables

List of Algorithms

Chapter 1

Introduction

The world's population is forecasted to increase from 7.8 billion in 2021 to 9.7 billion by 2050 (Worldometers.info, 2021). As such it is of critical importance to increase production of healthy and nutritious food through sustainable production methods. One potential source is seafood, which is rich in high quality protein, vitamins, minerals, and omega 3 acids. The world's consumption of seafood has increased from approximately 20 million tonnes in 1950 to 156 million tonnes in 2016, and is further projected to increase by 18 % by 2030 (FAO, 2020). As of 2018 aquaculture accounted for 52 % of consumed seafood, and its share is expected to rise to 59 % in 2030. Moreover, the amount of captured wild fish have remained flat during the last 30 years, and this trend is expected to continue until 2030. Meaning that a considerable part of the increase in seafood consumption must stem from increased aquaculture production. However, increasing aquaculture production is becoming increasingly difficult, as stricter environmental regulations are enforced world-wide. The need for regulation is illustrated by a disease outbreak which occurred in Chile during 2007 and 2008. The outbreak caused the country's salmon production tonnage to decrease by 64 % between 2008 and 2010, and production did not recover until 2012 as farmers were reluctant to deploy smolt in the aftermath of the outbreak (Vike, 2014).

Salmonids are a seafood type that accounted for 4.4 % of the global seafood consumption in 2018, and approximately 77 % of salmonids are farmed (Mowi ASA, 2020). Atlantic salmon makes up the largest share of salmonids on the market, and is becoming increasingly popular as production is up by a staggering 1000 % since 1990. Farmed Atlantic salmon are interesting from

a sustainability standpoint, as the carbon footprint per kilogram of edible salmonids are 65 % that of pork and 20 % that of beef, while the freshwater consumption per edible kilogram is 33 % and 13 % that of pork and beef respectively.

Biological and environmental constraints such as seawater temperature limit the potential locations of Atlantic salmon farms. As a consequence Atlantic salmon is only produced at sea in Norway, Chile, the United Kingdom, North America, the Faroe Islands, Ireland, New Zealand, and Tasmania. Norway is by far the largest producer of Atlantic salmon with a market share of 52 %, while Chile is the second largest producer with a market share of 28 %.

The Norwegian Atlantic salmon farming industry has since its infancy in the late 1960s increased its production significantly, through both technological advancements and increasing its production capacity. As the industry's ecological footprint has increased, more stringent regulations have been placed on the industry. One such regulation is that salmon farming companies must abide by biomass restrictions, which restrict how much biomass a company can have at sea at any point in time. These restrictions have caused the production volumes to stagnate as shown in Figure 1.1, and mean that a farming company can no longer increase its production by simply deploying more fish. Hence, the importance of utilizing biomass restrictions to the fullest have increased.

Another challenge is the ecological impacts salmon farming has on the environment, and the spread of diseases among farmed salmon has become a significant problem. A prevalent disease is salmon lice, which periodically appears at Norwegian salmon farms and causes both biological and economic damages. Meaning that farmers must account for the fact that salmon lice can disrupt their production schedules when they perform production planning.

In this thesis we investigate the potential of using novel smolt types in Atlantic salmon farming to try and increase production within the existing regulatory framework. Moreover, we model the presence and consequences of salmon lice on production planning through a stochastic model. We develop a multi-stage stochastic mixed integer program, and due to computational constraints we approximate this program through a two-stage program in

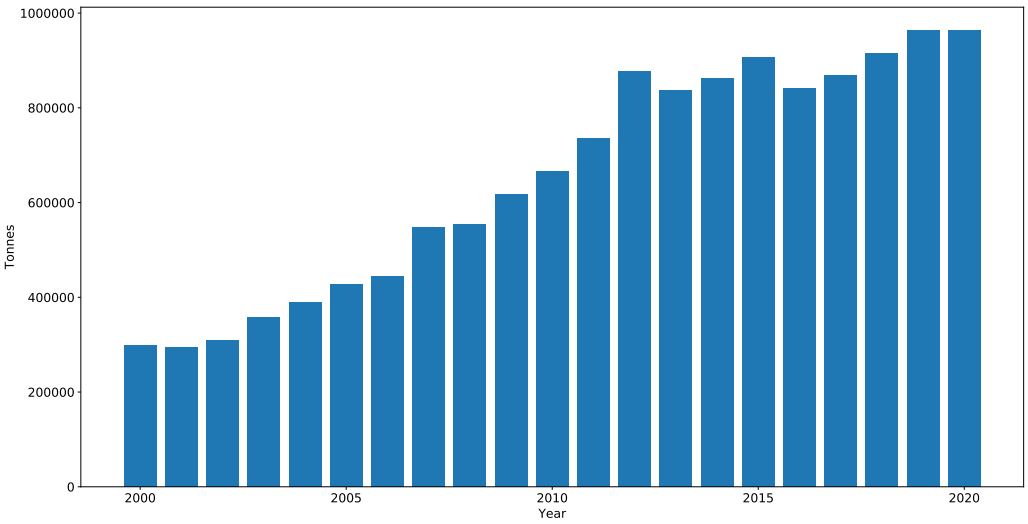

Figure 1.1: Yearly export of Norwegian farmed Atlantic salmon (Statistisk sentralbyrå, 2021).

combination with a simulation based rolling horizon heuristic framework to produce and evaluate operational plans.

This thesis consists of 11 chapters. We start off by providing some industry background on Atlantic salmon farming in chapter 2. Next, we present the problem description in chapter 3 that forms the basis for the rest of this thesis. We then review the relevant existing literature relating to the problem description in chapter 4. Chapter 5 goes through the basic modeling that is needed to explain the stochastic mixed integer linear program we introduce in chapter 6. We then discuss solution and evaluation methodologies in chapter 7. In chapter 8 we introduce a case study that we use to perform a computational study in chapter 9. We finish off by discussing some pathways for further research in chapter 10, before we give our concluding remarks in chapter 11.

Chapter 2

Industry Background

The Atlantic salmon farming industry's value chain can be broken up along the different stages of the salmon life cycle. Larger players usually integrate close to the entire value chain, while smaller players focus on specific parts of the value chain. As the salmon life cycle defines the value chain it is natural to start off by giving a short introduction to the life cycle as found in nature, before moving onto the industrial value chain. We then give a brief introduction to the nature of salmon prices. Next, we introduce the novel smolt types that are currently being studied by Aquagen AS. Thereafter we give an overview of causes of production losses including salmon lice. Lastly, we present relevant parts of the regulatory framework Norwegian Atlantic salmon farming companies must abide by.

2.1 Salmon Life Cycle

The salmon life cycle can be divided into seven stages, namely, eggs, alevins, fries, parrs, smolt, pre gender maturation salmon, and post gender maturation salmon. The stages preceding the smolt stage occur in freshwater, while the smolt and post smolt stages in general take place in saltwater. Salmon spawn as eggs in fresh water rivers usually during fall. The eggs hatch during the subsequent spring, and the salmon emerge as alevins. An alevin has its own yolk sac, which is its source of nutrition. After the yolk sac is consumed and the alevins have increased in size, they gain the ability to swim freely and turn into fries. Fries are extremely sensitive to environmental factors, thus the presence of fries in a river is a clear indication of a healthy aquatic

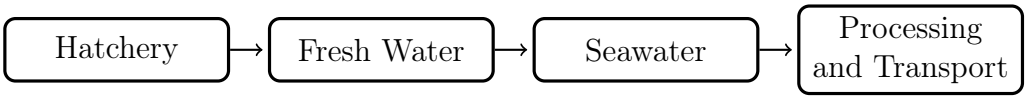

Figure 2.1: Main stages of the salmon farming value chain

environment. During the following fall fries develop vertical stripes and spots for the purpose of camouflage, and become classified as parrs. After approximately two to five years parrs undergo a transformation called smoltification to prepare for life in saltwater. The smoltification process turns fries into smolt, which migrate to seawater. When a salmon has grown sufficiently large it undergoes gender maturation and returns to fresh water rivers to breed. The precise timing of gender maturation is highly variable and genetically dependent.

2.2 Value Chain

The industrial farming value chain starts with the creation of eggs. Broodstock are used to breed salmon that have advantageous characteristics such as increased growth and reduced mortality rates. Breeding companies analyze the genes of their broodstock and select eggs and sperm from salmon containing suitable genetic profiles for spawning. Fertilized eggs are kept in a hatchery at approximately 8°C and are exposed to low light intensity (SinkabergHansen AS, 2021). During this phase unsuitable eggs are filtered out. Hatching usually occurs after approximately $400-500$ degree days, which in practice is close to $50-60$ days. In cases where the smolt producer differs from the egg supplier, eggs are often transported around 10 days before hatching to the smolt supplier.

After approximately $200-300$ degree days following hatching the fries' yolk sacs have been depleted by $2/3$, and the fries can now acquire nutrition from external sources (SinkabergHansen AS, 2021). Traditionally, this phase suffered from high mortality rates, however, modern techniques have achieved survivability rates close to 100 %. After approximately $8-10$ months smoltification occurs and the smolt is moved to seawater. The smoltification process can be sped up by using artificial lighting. The smolt can be kept in fresh water for a prolonged amount of time to increase its weight, in order to lower

mortality when it is released into seawater during months with lower temperatures. The smolt traditionally weigh between 100 – 250 gram at the time of seawater deployment.

Salmon is usually kept in seawater for approximately 12 – 18 months reaching a weight interval of 3 – 6 kg (Mowi ASA, 2020). During the seawater stage the salmon are exposed to environmental factors, and potentially salmon lice and other diseases. Disease outbreaks often necessitates either treatment or emergency harvesting. Both have economic consequences, as early harvesting means that the salmon will not reach is full potential weight, while treating diseases is costly in itself, but also disrupts growth rates and operational schedules. Lastly, at the end of the seawater stage salmon can undergo gender maturation which severely reduces meat quality and thus market prices. Using modern techniques such as artificial light one can postpone gender maturation enough to not be problematic (Iversen et al., 2016). However, research has shown that growth accelerates right before the occurrence of gender maturation, which could be exploited to increase production.

The final step of the seawater stage is to harvest the salmon. Together with the timing of deployment, timing of harvesting is one of the main parameters the farming company can control in order to maximize production. For welfare and meat quality reasons it is important that the salmon experience low stress levels during the harvest phase. To achieve this, the fish are starved for a minimum of one week before being anesthetized. After harvesting the salmon are transported to a slaughterhouse for gutting and further processing. Lastly, the processed salmon are transported to the end market. Mode of transportation varies, however, utilizing cargo space on civilian airfares is common practice.

2.3 Salmon Prices

An important aspect of salmon farming is the price the farming company can achieve for their salmon in the market. Salmon is priced per head on gutted (HOG) kilogram, however, the kilogram price varies according to how heavy the salmon is. Typically, larger salmon fetch higher kilogram prices than smaller salmon. Thus, one might be tempted to farm as heavy salmon as possible, however, this would come at the expense of increased feeding costs

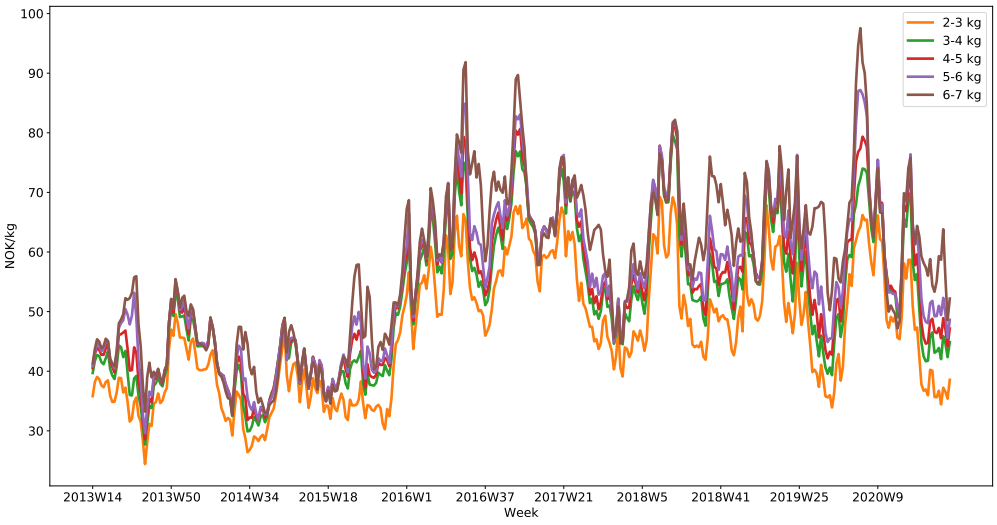

Figure 2.2: Historical weekly salmon prices per weight class as quoted on the NASDAQ Salmon Index.

and longer rearing cycles which ties up working capital. Moreover, salmon prices are highly volatile, which is evident from Figure 2.2, where weekly historical prices as quoted on the NASDAQ Salmon Index are presented for the most heavily traded weight classes. As the figure shows, the heavier weight classes consistently net higher prices than the lighter weight classes.

2.4 Smolt Types

We now give an introduction to the different smolt types that are currently studied by Aquagen AS. By introducing novel smolt types Aquagen aims to allow farming companies to increase their production tonnage. The current regulatory framework limits production tonnage indirectly through biomass restrictions. Thus, farming companies cannot simply deploy more salmon to increase production. They must instead utilize their biomass restrictions more efficiently, which is where the novel smolt types enter the picture.

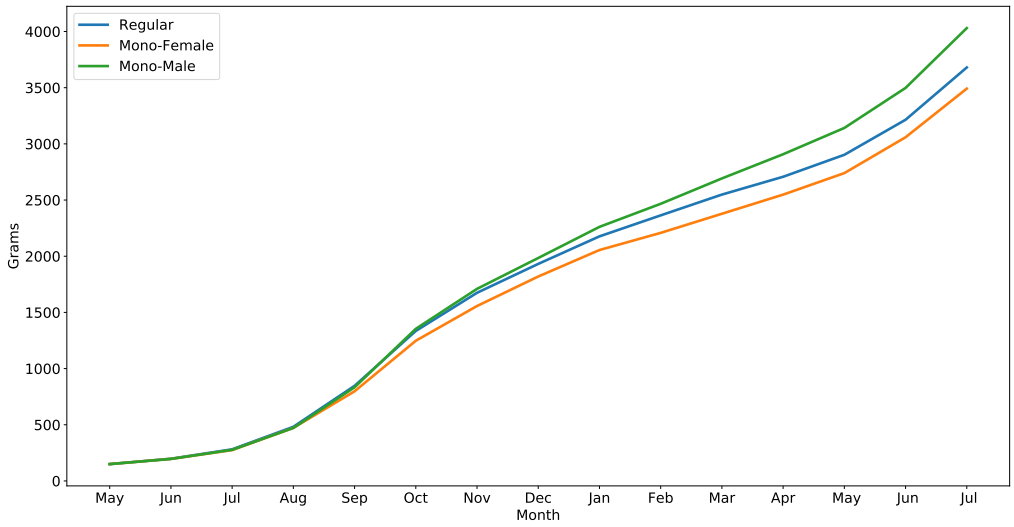

Figure 2.3: Weight development of a 150 gram smolt for both the mono-gender and regular smolt types based on growth rates obtained from sea trials.

2.4.1 Regular Smolt

The regular smolt type is the currently most typical smolt type that is used by Norwegian Atlantic salmon farmers. Although it is called regular, it is still a product of breeding. This thesis assumes that there only exist one type of regular smolt, however, in reality there exists several small variations that differ with regards to mortality and growth rates among other factors.

2.4.2 Mono-Gendered Smolt

Mono-gendered smolt refers to regular smolt where the two genders are kept separate from one another, e.g. a net-pen of only male smolt. Hence, mono-gender is not a feature of the smolt itself, but rather how it is reared. The genders are separated before entering seawater, as separating them during earlier stages are difficult as it is challenging to identify the different gender characteristics early on.

Sea trial results show that mono-male, mono-female, and regular smolt differ with respect to growth rates and mortality rates. The results reveal

that mono-female smolt have slower growth rates than the regular smolt, but also lower mortality rates. Moreover, mono-male smolt exhibited higher growth rates during the trials than both the mono-female and regular smolt. However, it also proved to have higher mortality rates than the two aforementioned smolt types. Figure 2.3 shows the weight development of a 150 gram smolt of the mono-gender and regular smolt types based on growth data obtained from sea trials.

2.4.3 Cooled Eggs

The cooled eggs smolt type has its name from the fact that the smolt hatch from eggs that are kept at lower temperatures than what has traditionally been used. Initial results from freshwater trials show that cooled eggs have a higher growth rate than a regular smolt type control group. Sea trials are currently ongoing, and there are hopes that the smolt originating from cooled eggs continue to exhibit higher growth rates.

2.4.4 Early Gender Maturation Smolt

Early gender maturation smolt are smolt that mature earlier than usual. This can be exploited as salmon experience increased growth rate in the time period preceding maturation. The hypothesis for the early maturation smolt type is that growth accelerates during the 13th – 14th month at sea. The disadvantage is that matured salmon has lower economical value due to its lowered meat quality. Therefore, it must be harvested before maturation takes place, limiting the amount of time the salmon can spend at sea to approximately 14 months. Preliminary results from ongoing sea trials indicate that growth does indeed increase in the time period leading up to maturation.

2.5 Production Loss

When farming Atlantic salmon production losses are common place. During a normal production cycle without any adverse events one would expect a total cumulative mortality rate of around 10 – 15 % (Bang Jensen et al., 2020). In this section we provide some background on some of the most common causes of production losses caused by adverse events.

2.5.1 Salmon Lice

Salmon louse is a copepod parasite living mainly on salmonids. The louse may damage the skin and fins of the salmon, and could potentially create skin wounds that open the salmon up for further infections from other sources. Salmon lice levels follow a seasonal pattern, mainly driven by fluctuations in seawater temperature. Salmon lice levels peak during summer and hit their low point during winter (Aldrin et al., 2013). Salmon lice are prominent in Norwegian Atlantic salmon farms and spread between farms if uncontained. From a farming standpoint one of the major complications of lice is the necessity to stop feeding while treating salmon lice, as the loss of feeding days reduces salmon growth rates. Additionally, mortality rates have a tendency to increase when treatments are performed.

Salmon lice also pose a significant threat to wild salmon, hence there is a clear need to limit the spread of salmon lice from farms. Thus, the Norwegian regulations stipulate that there must at all times be less than 0.5 female adult lice per fish in a salmon farm (Nærings- og fiskeridepartementet, 2012). However, during the weeks when wild salmon are migrating the threshold is lowered to 0.2. For Northern Norway the lowered threshold is applied during the weeks $21-26$, i.e. approximately late May through June. If the lice level rises above this threshold the farming company must take action to reduce the lice level, and possibly harvest the entire salmon population at the farm as an emergency solution. Several treatments for removing salmon lice exist. However, the lice tend to develop resistance against treatments. To avoid resistance problems, treatments techniques are rotated.

2.5.2 Other Diseases and Escapees

Other diseases also occur in Norwegian Atlantic salmon farms, such as the virus diseases pancreas disease (PD), infectious salmon anemia (ISA), infectious pancreatic necrosis (IPN), heart and skeletal muscle inflammation (HSMI), and cardiomyopathy syndrome (CMS) also commonly known as heart rupture (Fiskeridirektoratet, 2020). These diseases are often more serious than salmon lice infections, but occur less frequently. As an example, if the presence of ISA is proved, the affected net pens must be harvested in their entirety within a short time frame. This is in contrast to salmon lice, where treatment usually suffices to continue production without the need for emer-

gency harvesting. ISA infected salmon can safely be consumed by humans, however, some countries are hesitant to import infected fish. Moreover, if the salmon have been harvested at a weight below the marketable threshold they are more or less worthless. As diseases can cause significant economical damage, several methods such as vaccination and genetical breeding have been employed with success to eliminate several diseases from Norwegian Atlantic salmon farms.

2.6 Regulatory Framework in Norwegian Atlantic Salmon Farming

The regulatory framework for the Norwegian Atlantic salmon farming industry intends to both promote economical development of coastal regions, and fish welfare combined with a sustainable ecological footprint (Nærings- og fiskeridepartementet, 2008). In the following we give an introduction to the main regulatory requirements that farming companies must account for in their production plans.

2.6.1 Production Licenses and Biomass Regulations

Norwegian Atlantic salmon farming companies must abide by regulations that limit the amount of biomass they can have employed at sea at one point in time. The regulations regulate biomass both at a location, regional, and company level. A specific production location has a maximal allowable biomass (MAB) restriction assigned that limits the amount of biomass that can be employed at the location at any point in time. The Norwegian coast is divided into 13 production zones, and a company has its total employed biomass inside a production zone limited by the amount of production licenses she holds in the production zone. The regulations do however open for the possibility of combining licenses across production zones in some cases. A standard trout or Atlantic salmon production license has a maximum allowable biomass of 780 tonnes outside of the northernmost county of Troms and Finnmark where it is increased to 945 tonnes.

Figure 2.4 gives a graphical explanation of the MAB system for a Norwegian Atlantic salmon farming company. In zone 1 the company owns 8 licenses of 780 tonnes each, limiting the total employed biomass in the zone to 6,240

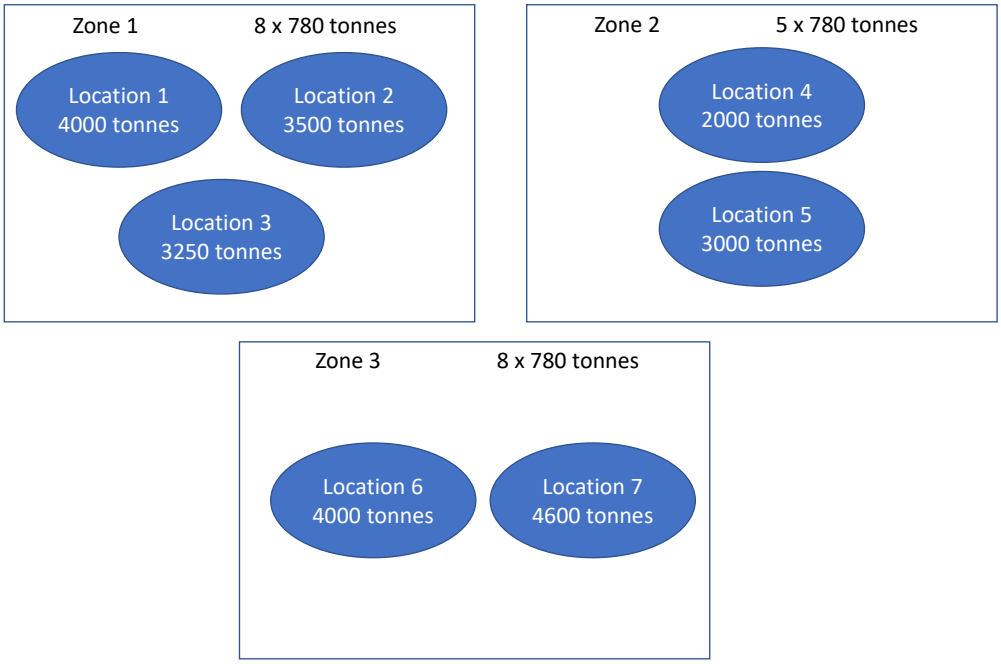

Figure 2.4: Graphical explanation of the regulatory MAB system at a regional and location level.

tonnes. Besides the regional MAB, the company must also respect the location MAB, meaning that the company can never have more than 4,000 tonnes employed at location 1, and can similarly not exceed 3,500 tonnes and 3,250 tonnes at location 2 and 3 respectively. In the case where the company has been granted an inter-regional MAB between two or more zones, the company is allowed to sum all licenses in the relevant zones, and can in reality view the zones in question as merged into one zone. Assume that the company depicted in the figure has an inter-regional MAB between zone 1 and 2, then the regions are in reality merged with a total of $8 + 5 = 13$ licenses and a combined regional MAB of $(8+5) \cdot 780 = 10,140$ tonnes. In the extreme case, the company could potentially have 10,140 tonnes employed in zone 1 and 0 tonnes employed in zone 2. In some cases, such as the case study we introduce in chapter 8, the company is granted an inter-regional MAB between all of her zones, yielding a company wide MAB that replaces several regional MABs, however, as always the location MABs must still be

respected.

To counteract stagnating production volumes caused by a lack of new production licenses, the authorities introduced the so called traffic light system in 2017 (Nærings- og fiskeridepartementet, 2017). The goal of the traffic light system is to allow new licenses to be issued in an ecologically safe manner. In short, each of the 13 production zones are assigned one of the three colors red, green, and yellow, hence the name traffic light system. A zone that is red must reduce its total allowable biomass by 6 %, yellow zones maintains the status quo, while green zones can increase its total allowable biomass by 6 %. The zones' colors are decided by the impact of salmon lice on wild smolt mortality. If less than 10 % of the wild smolt in the production zone are expected to die as a consequence of salmon lice the zone is colored green, mortality in the range $10-30$ % colors the zone yellow, and if it is above 30 % the zone is colored red. The coloring as of April 2021 is given in Figure 2.5.

2.6.2 Operational Requirements

The Norwegian Atlantic salmon farming regulations require that companies every year produce a two-year operational plan. This plan must contain information on where, when, and by which amount the company plans to deploy smolt, and similarly where and when she plans to harvest. Moreover, the company must ensure that she fallows a location for a minimum of two months between completing harvesting of a deployment and deploying smolt again. If a location remains unused for 24 months or longer the company risks having the production license for that location withdrawn.

As we touched upon in section 2.5, farming companies are required to monitor and report the presence of infectious diseases to the authorities, and potentially act through treating or emergency harvesting the salmon.

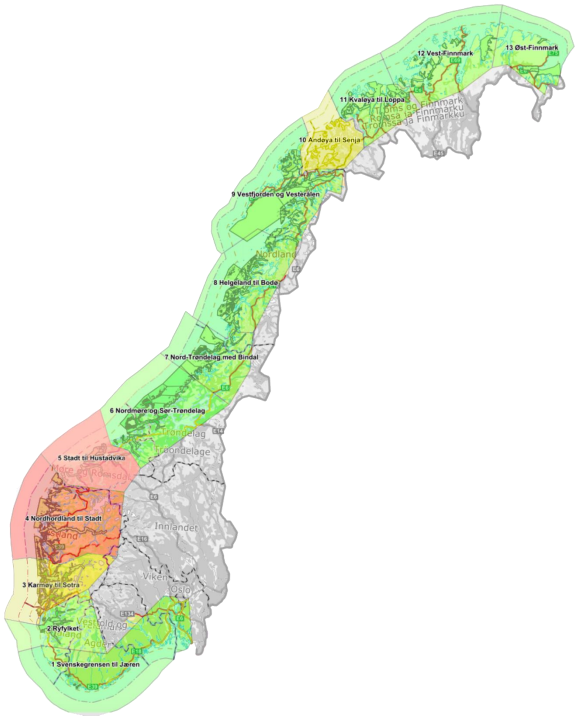

Figure 2.5: Status of the traffic light system for each of the 13 zones as of April 2021.

Chapter 3

Problem Description

In this chapter we describe the stochastic tactical optimization problem an Atlantic salmon farming company faces during the seawater stage of the value chain. Thus, in accordance with Anthony's framework (Anthony, 1965), the scope of the problem is limited to aspects affecting deployment of smolt and harvesting of salmon over a mid-term planning horizon.

The company must decide when, where, and by which amount she wishes to deploy smolt. She also needs to decide which smolt types she wishes to deploy, as different types inhibit different mortality and growth characteristics. Moreover, the company must also decide the initial weight of the smolt that is to be deployed. Both the smolt type and initial weight must be chosen from seasonally dependent sets of available smolt types and weights. A deployment of smolt with the same smolt type and weight specification at a location constitutes a cohort. The company can decide to deploy one or more cohorts at the same time at a location. Moreover, during certain time periods the supply of smolt may be limited.

Operational limitations and practical considerations limit the amount of fish that can be deployed as part of a cohort at a location to lie between an upper and lower bound during a time period. Moreover, biological considerations necessitate that deployment of smolt is limited to certain periods of the year. When smolt has been deployed at a location, no other deployments can take place at the same location during the rearing cycles of the deployed cohorts.

After deploying a cohort, the company must decide when to harvest it.

The mean weight of a deployed cohort must have reached a specified weight threshold before it can be harvested, and must not have surpassed an upper weight threshold. However, the company need not harvest the entire cohort at the same time. The company must limit her harvest volumes at location during a time period to lie between a lower and upper bound if she chooses to harvest. Additionally, her company wide harvest volumes during a time period are restricted by a capacity limit. Moreover, the company restricts the maximum length of a rearing cycle, requiring that a deployed cohort must be harvested before its time at sea surpasses this restriction. After harvesting the company sells her harvested salmon into a market that prices salmon based on different weight classes. The price within each weight class fluctuates over time.

The company has several locations, each of which has a given maximum allowed biomass, that restricts the total amount of biomass that can be employed at the location at any point in time. Additionally, a company wide maximum allowed biomass restriction, limits the sum of simultaneously employed biomass across all locations. The company must also ensure that a location is fallowed for a minimum amount of time after a production cycle. However, a location must never be fallowed for more than a maximum amount of time, otherwise its location license could be withdrawn.

The company bases her tactical decisions on growth and mortality predictions. However, both growth and mortality rates can be impacted by salmon lice treatments. The necessity of salmon lice treatments is stochastic and the company has several scenarios for how the necessity of salmon lice treatments will impact biomass development during a time period at a location level.

The time periods that constitute the company's planning horizon is divided into multiple disjoint sets of contiguous time periods called planning blocks. For each planning block the company must decide her deployments together with where and when to fallow and harvest before the planning block starts, and she cannot change these decisions after the planning block has started. However, she need not decide the harvest volumes before the planning block commences, she can instead wait until new information arrives, revealing the growth and mortality rates her deployed cohorts has experienced.

Finally, the company must ensure that there is a certain amount of biomass

employed at the end of the planning horizon, and that this biomass will not violate maximum allowed biomass regulations in the time following the planning horizon.

Chapter 4

Literature Review

In this chapter we review existing literature relating to the problem description given in chapter 3. We start off by examining the literature on general operations research approaches to aquaculture and specifically salmon farming. Thereafter, we review relevant literature regarding the biological processes of biomass development, mortality, and salmon lice. Next, we give an overview of the relevant literature on salmon prices. Lastly, we outline the contribution of this thesis.

4.1 Production Planning

For a general overview of existing literature on operations research applied to aquaculture we refer the reader to Bjørndal et al. (2004), which gives a broad overview of different historical approaches to aquaculture operations optimization.

Bjørndal (1988) develops a bioeconomic model to determine the optimal timing of harvesting a cohort. The optimal time for harvesting is in general when marginal revenue equals marginal costs, which is a well known microeconomic axiom. He studies how the optimal timing of harvesting changes when different costs are accounted for. He finds that accounting for a harvesting cost per fish pushes the time point for optimal harvesting out in time, while introducing feeding costs brings it forward in time. Moreover, he concludes that a graded harvesting scheme is superior to a batch harvesting scheme. In a graded harvesting scheme the farming company can decide which weight

classes to harvest from in a cohort during a harvest operation. This is in contrast to a batch harvesting strategy where the company harvests across the entire cohort without filtering for certain weight classes.

Arnason (1992) investigates the interplay between feeding and timing of harvesting, which he finds to be heavily interdependent. In addition, the optimal feeding schedule depends on the applied growth model, and there are few generalizations that hold across different growth models. Heaps (1993) extends the work of Arnason (1992) by examining the potential in letting the feeding rate be a decision variable controlled by the farming company, under the assumption of a density independent growth model. A consequence of having the ability to control the feeding rate is that the company can control the growth rate, instead of having the growth rate dictate the feeding rate. He finds that culling is not always a part of an optimal operational schedule. Culling is the practice of performing graded harvesting at different points in time during a rearing cycle. A density dependent growth model is introduced in Heaps (1995), and results regarding feeding schedules are shown to generalize from the density independent growth models introduced in Heaps (1993). However, with a density dependent growth model, culling might be part of an optimal operational plan.

Forsberg (1996) presents a multi-period linear program that determines optimal deployment and harvesting volumes for a single fish farm with a total allowable biomass constraint. The model optimizes discounted profits and accounts for weight dependent price levels on the revenue side, while incorporating smolt, feeding, and harvesting costs on the cost side. Biomass development is modeled using discrete weight classes and a time varying Markov process describing transition probabilities between weight classes. In Forsberg (1999) he develops a slightly more comprehensive multi-period linear program, and determines that pursuing a graded harvesting strategy can increase profits by 11.8 % compared to performing batch harvesting.

Pascoe et al. (2002) argue that existing optimization models for aquaculture do not properly capture uncertainty, and as a consequence their output cycle lengths deviates significantly from real world observations. They point out that discount rates are neither continuous nor constant during rearing cycles, which is assumed in most models. In a case study of tiger prawns they find it is necessary to set the discount rate to a staggering 400 % to

approximate the rearing cycle lengths observed in industry.

Guttormsen (2008) develops an extended version of the Faustmann model introduced in (Faustmann, 1849) to solve the rotational problem in fish farming. He finds that the optimal cycle length is 19 months in the simplest application of the model to Atlantic salmon farming. In further experiments he accounts for both the cyclical nature of relative prices, and climate restrictions limiting deployment of different fish types to specific parts of the year. Results show that accounting for these two factors affect the optimal cycle length both in isolation and in the joint case. However, whether it increases or decreases the optimal cycle length depends on which calendar month deployment takes place.

Hæreid et al. (2013) take a multi-stage stochastic programming approach to maximizing expected profits for a production planning problem spanning $4-5$ years for an Atlantic salmon farming company with multiple locations. They account for uncertainty in prices, growth, and mortality. They base their biomass development on a growth table provided by the feeding company Skretting AS, and apply interpolation to increase the table's granularity. Bravo et al. (2013) construct two mixed integer linear programs (MILP), one for the freshwater stage of the value chain, which aims to minimize costs while satisfying the seawater stage's demand, and the second for the seawater stage with the goal to maximize harvested biomass while satisfying demand from a processing plant. Their models are aimed at giving decision support to one of Chile's largest salmon farming companies. Preliminary results show that the models help the farming company abide by constraints that are difficult to enforce using manual planning, and in doing so also increased their objective values.

Næss and Patricksson (2019) create a two-stage stochastic MILP, which maximizes the expected total harvest volumes over a time horizon of five years. They apply the model with a monthly time resolution on a farming network consisting of 16 farms. The model incorporates regulatory requirements such as maximum allowable biomass constraints, and activity and fallowing requirements. The model is thus able to create production plans that utilize the full flexibility the regulatory framework offers, instead of pre-defining certain decisions to abide by the regulations. However, this also increases the computational complexity, they thus aggregate the 16 locations into three

locations to solve stochastic instances of five scenarios, and apply sample average approximation. They also study the increase in harvest volumes that can be achieved by introducing male mono-gender smolt. They model mono-gender growth rates by scaling the growth table used in Hæreid et al. (2013) based on data from sea trials. Their deterministic study finds that introducing male mono-gender smolt could increase production tonnage by 21.3 %, and that male mono-gender smolt is used exclusively when available. While, their stochastic studies show that having stochastic temperatures reduces the total expected harvest volume by 4.2 % compared to the deterministic study.

Oglend and Soini (2020) study the effect of having a regulator limiting supply of aquaculture licenses. In general they find that limiting maximum allowable biomass constraints result in companies increasing stocking densities when optimizing for profits, meaning the farming companies harvest at lower weights than they would have if they had unconstrained production capacity.

4.2 Biological Processes

We now outline the literature regarding the relevant biological processes.

4.2.1 Biomass Growth

Growth is one of the most important factors to account for when making operational decisions in Atlantic salmon farming. As a consequence several growth models of various complexity and robustness have been developed. We refer the reader to Aunsmo et al. (2014) for a comprehensive review of growth models for Atlantic Salmon.

The first growth model we review is the average daily growth rate (ADG), which is given by

$$\text{ADG} = \frac{w_n - w_1}{\sum_{d=1}^{n} d}, \tag{4.1}$$

where w_n and w_1 denote the weight of the fish at day n and 1 respectively. Aunsmo et al. (2014) find ADG to be inappropriate for fish growth modeling, as absolute weight growth increases with the weight of the fish. Thus, it is sensitive to the initial stocking weight of the fish.

The daily specific growth rate (SGR) introduced by Brett and Groves (1979) can be calculated using the same notation as with ADG as

$$\text{SGR} = \frac{\ln(w_n) - \ln(w_1)}{\sum_{d=1}^{n} d}. \tag{4.2}$$

Forsberg (1995) approximates the relationship between SGR, the seawater temperature T in Celsius, and fish weight w in grams as

$$\text{SGR} \approx \text{T}^{0.97} w^{-0.34}. \tag{4.3}$$

This relationship is limited to Atlantic salmon with a weight between 50 – 3000 grams, and seawater temperatures in the range 4°C – 14°C. As equation (4.3) shows, SGR varies with the weight of the fish. This pattern is also found by Aunsmo et al. (2014), as they find stocking weight to influence SGR.

Aunsmo et al. (2014) find the thermal growth coefficient (TGC) by Iwama and Tautz (1981) to be the most suitable model among the non-proprietary models they examine. It is defined as

$$\text{TGC} = 1000 \cdot \frac{w_n^{\frac{1}{3}} - w_1^{\frac{1}{3}}}{\sum_{d=1}^{n} \text{T}_d}, \tag{4.4}$$

where T_d is the average seawater temperature measured in Celsius during day d. Jobling (2003) cautions against uncritical use of TGC across different temperatures, and shows how it deviates outside of normal temperature ranges for different fish types. Moreover, Aunsmo et al. (2014) find TGC to depend on stocking weight, but to a lesser degree than ADG and SGR. Thorarensen and Farrell (2011) state that the difference in TGC for a fish with a specific weight varies less than 4 % in the thermal range between 4°C and 14°C. Moreover, the weight difference for a fish weighing 50 grams at 4°C and a fish weighing 3000 grams at 14°C is less than 6 %.

The Ewos growth index (EGI) is a proprietary regression based growth model. Letting α be a vector of regression parameters including an intercept, and letting $e_i \sim \mathcal{N}(0, 1)$ be an error term we have

$$\text{Y}_i^{\text{EGI}} = \alpha_0 + \alpha_1 \text{T}_i^2 + \alpha_2 \text{T}_i^4 + \alpha_3 x_{1i} + \alpha_4 x_{2i} + \alpha_5 w_i^{\frac{1}{3}} + \alpha_6 x_{1i} w_i^{\frac{1}{3}} + \alpha_7 w_i^{\frac{2}{3}} + e_i, \tag{4.5}$$

where T_i is the seawater temperature, x_{1i} is the length of the day, which is defined as the time the sun spends above the horizon, x_{2i} is the change in

day length, and w_i is the fish weight. Aunsmo et al. (2014) find EGI to be the most suitable model overall, however, as mentioned it is proprietary and the regression parameters are thus not publicly available.

4.2.2 Mortality Patterns

Bang Jensen et al. (2020) examine mortality for Norwegian Atlantic salmon farms during the years 2014 – 2018. They define the mortality rate during a time period t, Ψ_t, as

$$\Psi_t = \frac{d_t}{\frac{1}{2}(n_{t+1} + n_t)}, \tag{4.6}$$

where d_t is the number of fish that died during time period t, and n_t is the number of fish alive at the start of time period t. The denominator represents the average amount of fish alive during time period t, under the assumption of a linear reduction in the cohort's population. Moreover, they define the mortality risk μ_t, which is the probability of an individual fish dying during time period t, as

$$\mu_t = 1 - e^{-\Psi_t}. \tag{4.7}$$

They estimate the cumulative mortality risk for a cohort that has spent 15 months at sea to lie in the interval 10 – 15 % on average. Moreover, they investigate the connection between months spent at sea and monthly mortality risk. They find the mortality risk to be at its highest during a cohort's first month at sea, thereafter it reaches its minimum during the third month at sea, before it starts increasing again and plateaus after approximately 14 months. A graphical replication of their findings is shown in Figure 4.1. A similar pattern is established in Scotland by Salama et al. (2016) during the years 2002 – 2009.

4.2.3 Salmon Lice Prediction

Salmon lice is one of the main causes of production loss in the Norwegian Atlantic salmon farming industry, and thus predicting the presence of salmon lice has seen extensive study. Aldrin et al. (2013) develop a stochastic space-time model where the monthly lice abundance level for every Norwegian Atlantic salmon farm is predicted simultanously. They find that 66 % of expected lice abundance at a farm is explained by internal factors at the farm, while 22 % is attributable to neighboring farms, and the final 6 % to

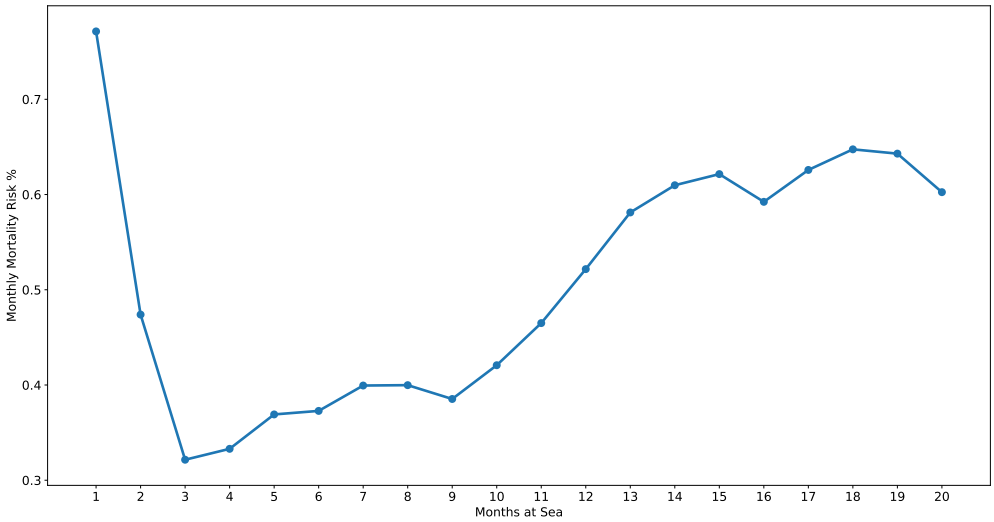

Figure 4.1: Monthly mortality risk of a cohort as a function of months spent at sea in Norway during the years 2014 through 2018. Replicated from Bang Jensen et al. (2020).

non-specified factors. They model the correlation between farms as exponentially decreasing with seaway distance. Kristoffersen et al. (2014) build on Aldrin et al. (2013) and take a biological approach in modeling the lice infection pressure by accounting for egg production and different stages of the louse life-cycle. Aldrin et al. (2017) expand further on the aforementioned works by introducing a Bayesian mechanistic population model at individual farms, where the parameters are obtained from full scale production data.

Overton et al. (2019) survey salmon lice treatment methodologies in Norway during the years 2012 – 2017, and find that treatment methodologies varies over time and between production zones. They find a clear shift from the use of chemotherapeutant baths during 2012 – 2015 to thermal, mechanical, and hydroperoxide treatments during 2016 – 2017. Moreover, they find an increase in mortality during months where lice treatments are performed, and that the size of the increase depends on temperature, fish size, and the cohorts' previously experienced mortality rates.

Kragesteen et al. (2019) build a bioeconomic model and study the lice level

threshold a company should use to maximize profits. The lice level threshold is used to indicate when salmon lice treatment should commence at a farm. They show that when optimizing for an individual farm the company should use a rather high threshold. However, when optimizing for an entire network of farms a threshold of 0.1 gravid female lice per salmon should be used, which is considerably lower than the Norwegian government's requirement of 0.5. They conclude that the difference in the optimal threshold between optimizing in isolation and in a network gives rise to a tragedy-of-the-commons environment. Meaning, that from a game theoretic standpoint, strict governmental enforcement is needed to ensure that the different farmers abide by the common optimal threshold.

4.3 Patterns in Salmon Prices

Multiple studies find that salmon prices are non-stationary and do not exhibit persistent seasonal patterns (Asche et al., 1999; Asche and Guttormsen, 2001). Asche and Guttormsen (2001) study seasonal patterns of relative differences in price between different sizes of salmon. They define the relative price P_t^α per kilogram of weight class α at time t as

$$\mathrm{P}_t^\alpha = \frac{\hat{\mathrm{P}}_t^\alpha}{\hat{\mathrm{P}}_t^{\bar{\alpha}}}, \tag{4.8}$$

where $\bar{\alpha}$ is the reference weight class, and $\hat{\mathrm{P}}$ denotes the absolute price per kilogram. They study the prices of weight classes 1 – 2 kg, 2 – 3 kg, 3 – 5 kg, 5 – 6 kg, and 6 – 7kg, and they use the aggregated 3 – 5 kg weight class as the reference weight class. They find that the relative price of a weight class is cyclical with a cycle length of approximately one year, and that this seasonal pattern can be approximated by trigonometric regression.

In general, the relative prices of larger salmon weighing 5 – 7 kg reach their high point in August and September, where the 6 – 7 kg class has a price of approximately 120 % relative to the 3 – 5 kg reference class, and their relative prices reach their low point at around 90 – 95 % in February and March. Smaller salmon weighing 1 – 3 kg exhibit their lowest relative price during May through June, and their relative prices are at their highest during November through January. They hypothesize that the seasonal patterns

are caused by the timing of traditional rearing cycles. Typically, companies prefer to harvest large salmon before August as the likelihood of gender maturation is at its peak during August and September, thus lower volumes of large salmon are harvested during these two months. Moreover, they find that both the absolute and first differences in price of neighboring weight classes, e.g. $1-2$ kg and $2-3$ kg, are highly correlated, while the correlation between $1-2$ kg and $6-7$ kg is much lower.

Forsberg and Guttormsen (2006) examine both absolute and relative salmon prices. They show that absolute prices tend to peek during weeks $20-24$ of the year, i.e. the mid of May to the mid of June, while they bottom out during November. The difference between the highest and lowest prices achieved during a year is usually around 20 %. Interestingly, they find that absolute prices do not show any significant increases during Christmas, which is commonly believed in the industry. This indicates that salmon farmers adapt to the high demand around Christmas by increasing supply. They also find the seasonal patterns in relative prices found in Asche and Guttormsen (2001) to still be present. Moreover, they study the value of information that prices inhibit by applying the model introduced in Forsberg (1999). They find that the value of information of prices is significant, as a farmer who plans based on perfect price information for each weight class could triple his profits compared to one who plans according to historical data without separating weight classes. They also find that planing based on historical data for each weight class instead of prices that aggregate the weight classes can increase profits by approximately 33 %.

Asche et al. (2016) examine the relationship between futures prices and spot prices for salmon on the Fish Pool exchange. They consider futures contracts with up to six months maturity, and find the futures price to provide an unbiased estimate for the spot price. However, this relationship is unidirectional. Thus, futures prices does not provide a price discovery function, and they therefore conclude that the salmon futures market is still immature.

4.4 Our Contributions

We develop a multi-stage stochastic mixed integer linear program, which is inspired by Næss and Patricksson (2019), that accounts for regulatory

and practical constraints. Biomass development is modeled according to the thermal growth coefficient and accounts for patterns in mortality. We replace the discrete weight classes and transition probabilities Næss and Patricksson (2019) and Hæreid et al. (2013) use to model biomass development, with a cohort-level biomass development scheme, which reduces solution-time significantly. Moreover, we model the probability of lice treatments at a salmon farm with dependence between the farms in the farming network, and the interdependence between lice treatments and growth and mortality rates. Finally, we introduce a rolling horizon framework to evaluate solutions of a simplified two-stage version of the multi-stage program in a computationally tractable manner. We study both the potential of introducing novel smolt types, and the difference between maximizing raw harvest volumes and maximizing with respect to relative prices. Lastly, we study the impact of using stochastic programming when planning versus planning deterministically assuming that no lice treatments are necessary.

Chapter 5

Methodology

In this chapter we go through the modeling that is needed to develop the mathematical model that we introduce in chapter 6. We start off by introducing a model for biomass development, next we describe the stochastic process of salmon lice presence and its connection to biomass development. We then develop a framework for weight distribution of a cohort. Lastly, we examine updated relative price data.

5.1 Biomass Development

Biomass growth is the most pivotal parameter to model in the salmon farming industry. Several growth models exist, and they have different strengths and weaknesses. We choose to use the thermal growth coefficient (TGC), which is a simple model that have been found to be robust, and does not rely on any proprietary parameters or data. Using TGC we can predict the development of the mean weight of a cohort according to

$$w_{t+1} = (w_t^{\frac{1}{3}} + \frac{1}{1000} \cdot \mathrm{TGC}_t \cdot \mathrm{L}_t \cdot \mathrm{T}_t)^3, \tag{5.1}$$

where w_t is the mean weight of the cohort's salmon at the start of time period t, L_t is the length of time period t in days, and T_t is the mean daily seawater temperature during the time period in Celsius. We have introduced a time index for the TGC parameter, which is intended to adjust for the fact that TGC depends on the amount of time the fish have enjoyed at sea.

The salmon also experience a mortality rate μ_t during time period t. Letting the fraction of fish still alive at time t be s_t, we have the recursive relationship

$$s_t = (1 - \mu_{t-1})s_{t-1}. \tag{5.2}$$

At the time of deployment the entire cohort is still alive, giving $s_1 = 1$. We can then define a parameter A_t as the ratio between a cohort's total biomass at the start of time period t and its initial total biomass

$$\mathrm{A}_t = \frac{s_t \cdot w_t}{w_1}. \tag{5.3}$$

Assuming that no harvesting takes place, the total biomass of a cohort at the start of time period t, which we denote M_t, is simply given by

$$\mathrm{M}_t = \mathrm{A}_t \cdot y, \tag{5.4}$$

where y is the initially deployed biomass. However, when harvesting is performed we need to adjust for the harvested volumes. For the sake of this thesis we assume that harvesting is performed according to a batch harvesting strategy. We define the parameter R_t to be the percentage increase in total biomass during time period $t-1$, i.e. between the start of time period $t-1$ and and the start of time period t, assuming no harvesting have taken place

$$\mathrm{R}_t = \frac{\mathrm{A}_t}{\mathrm{A}_{t-1}}. \tag{5.5}$$

Denoting the total harvest volume during time period t as H_t, we assume that the following relationship holds

$$\mathrm{M}_t = \mathrm{R}_t \cdot (\mathrm{M}_{t-1} - \mathrm{H}_{t-1}). \tag{5.6}$$

We note that although this thesis specifically uses TGC, the way biomass development is incorporated in the stochastic mixed integer linear program we introduce in chapter 6 is independent of the underlying biomass development model, it suffices to provide a specification of A_t and thus implicitly R_t.

5.2 Salmon Louse Modeling

The salmon lice level at a salmon farm is a stochastic process. When lice levels pass a certain threshold treatment is necessary, both from a regulatory

perspective and from the perspective of protecting the salmon against getting eaten by lice. In order to treat lice the salmon farmer must stop feeding the fish, and the fish might require time to recover before they resume acquiring nutrition after the treatment have taken place. This means that growth rates slow during lice treatments. We start off by modeling the presence of lice treatments at an individual location in isolation, before moving onto modeling dependence between locations in a farming network, and finish off by describing salmon lice treatments' impact on biomass development

5.2.1 Individual Location

We let ρ_{it} be a random variable that denotes the probability that treatment of salmon lice is necessary during time period t at location i. Moreover, we let ζ_{it} be a binary random variable representing whether a lice treatment is necessary during time period t at location i that is distributed according to

$$\zeta_{it} \sim \text{Bernoulli}(\rho_{it}). \tag{5.7}$$

We assume the probability of a lice treatment being necessary during time period t depends on two factors, namely, calendar month and whether treatment took place during the previous time period. Note that we omit the aspect of whether biomass is present at a location to reduce the dimension of the state-space. For notational ease we introduce $\hat{\rho}_{it} = \text{P}(\zeta_{it} = 1|\zeta_{i(t-1)} = 0)$, i.e. the probability of a treatment at time t at location i if no treatment was performed at time $t-1$ at location i, and $\check{\rho}_{it} = \text{P}(\zeta_{it} = 1|\zeta_{i(t-1)} = 1)$ representing the probability of treatment given treatment took place during the previous time period. We then assume the following relationship

$$\rho_{it} = \hat{\rho}_{it}(1 - \zeta_{i(t-1)}) + \check{\rho}_{it}\zeta_{i(t-1)}, \tag{5.8}$$

where we let ζ_{i0} be a constant present through initial conditions.

Figure 5.1 depicts empirical measurements of ρ_{it} in blue, $\hat{\rho}_{it}$ in orange, and $\check{\rho}_{it}$ in green, based on data from Barentswatch (2021) for production zone 9 – Vestfjorden and Vesterålen in Northern Norway for the years 2012 – 2020. When estimating these probabilities we have only considered farms that are denoted as likely active during both time period t and time period $t-1$. Moreover, we have assumed that the probabilities are equal for all farms located in the production zone.

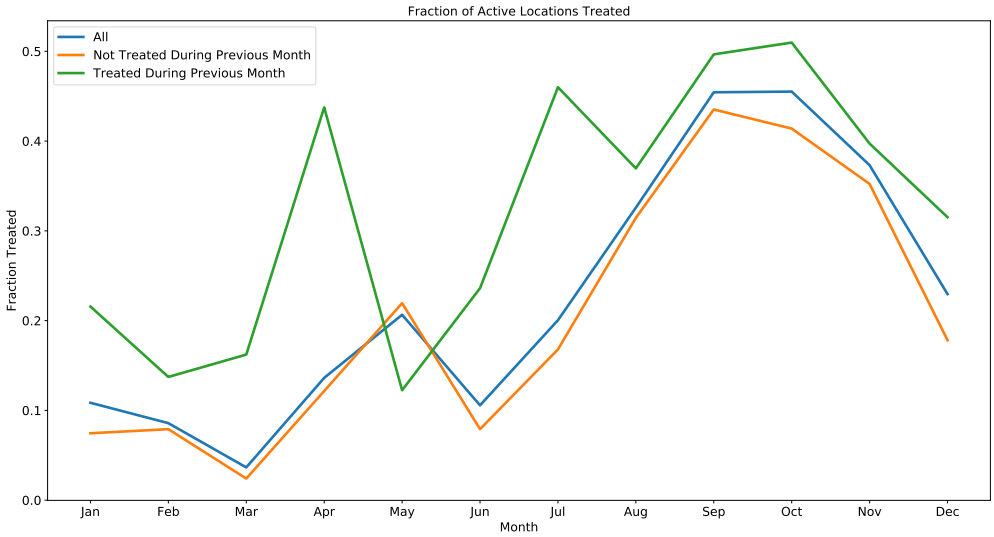

Figure 5.1: Fraction of active locations in production zone 9 treated for lice during different calendar months.

5.2.2 Correlation Between Locations

If a treatment is necessary at a specific location, there is naturally an increased probability that treatments are necessary at nearby locations. Hence, given two locations i and j we have dependence between ζ_{it} and ζ_{jt}, which is attributable to the spatial dimension. We express this dependence through the relative risk of infection, RR. Letting RR_{ij} denote the relative risk of infection between location i and j, we have the following definition

$$\mathrm{RR}_{ij} = \frac{\mathrm{P}(\zeta_{it} = 1 | \zeta_{jt} = 1)}{\mathrm{P}(\zeta_{it} = 1 | \zeta_{jt} = 0)}. \tag{5.9}$$

Aldrin et al. (2013) gives the relation between RR_{ij} and the seaway distance d_{ij} between location i and j as

$$\mathrm{RR}_{ij} \propto e^{-\frac{\phi_1}{\phi_2} d_{ij}^{\phi_2}}, \tag{5.10}$$

where ϕ_1 and ϕ_2 are constant parameters. As the presence of salmon lice at another farm cannot reduce the risk of infection, we have that the relative risk is not less than one in our case. We thus assume the following

$$\mathrm{RR}_{ij} = \max\{\mathrm{N} \cdot e^{-\frac{\phi_1}{\phi_2} d_{ij}^{\phi_2}}, 1\}, \tag{5.11}$$

where N is a constant parameter deciding at which distance the relative risk becomes one. Combining equation 5.9 with the law of total probability, we arrive at the joint probability distribution between two pairwise locations

$$\mathrm{P}_{ij}^{t} = \mathrm{P}(\zeta_{it} = 1 \wedge \zeta_{jt} = 1) = \frac{\rho_{it}\rho_{jt}}{1 + \rho_{jt}(\mathrm{RR}_{ij} - 1)}\mathrm{RR}_{ij}. \tag{5.12}$$

As the reader may note, it is possible that $\mathrm{P}_{ij}^{t} \neq \mathrm{P}_{ji}^{t}$ when $\rho_{it} \neq \rho_{jt}$. To correct for this issue we introduce a corrected joint probability $\mathrm{P}_{ij}^{\prime t}$, which must also respect the property $\rho_{it} \geq \mathrm{P}_{ij}^{\prime t}$, where the latter requirement follows directly from the law of total probability. A natural choice thus becomes

$$\mathrm{P}_{ij}^{\prime t} = \min\{\mathrm{P}_{ij}^{t}, \mathrm{P}_{ji}^{t}\}. \tag{5.13}$$

We can now estimate the joint probability distribution of a treatment being necessary $\mathrm{P}(\bigwedge_{i\in\mathcal{I}}(\zeta_{it} = c_{it}))$ for all locations, by utilizing iterative proportional fitting (IPF). We refer the reader to Gange (1995) for a detailed discussion of IPF. For the sake of our applications IPF estimates a joint probability distribution respecting both the marginal distribution for each location and the marginal joint distribution between pairwise locations. In practice, we only estimate the marginal distribution for each location, and the Pearson correlation coefficient between all pairwise locations, the rest is handled by the R procedure RMultBinary in the mipfp package. For completeness we state the Pearson correlation coefficient for ζ_{it} and ζ_{jt}

$$\mathrm{corr}(\zeta_{it}, \zeta_{jt}) = \begin{cases} 1 & i = j \\ \frac{\mathrm{P}_{ij}^{\prime t} - \rho_{it}\rho_{jt}}{\sqrt{\rho_{it}\rho_{jt}(1-\rho_{it})(1-\rho_{jt})}} & i \neq j \wedge \rho_{it} > 0 \wedge \rho_{jt} > 0 \\ 0 & otherwise \end{cases}. \tag{5.14}$$

5.2.3 Impact of Salmon Lice on Biomass Development

The main point of interest in modeling the presence of lice treatments is to capture its effect on growth and mortality rates. For our purposes lice treatments affect growth rates through a reduction in feeding days. We let the amount of feeding days during time period t, L_{it}, be a random variable defined as

$$\mathrm{L}_{it} = \mathrm{L}_{t} - \mathrm{D} \cdot \zeta_{it}, \tag{5.15}$$

where L_t is the amount of days in time period t and D is the amount of feeding days lost during a treatment. We assume that a maximum of one

treatment can take place per time period. The change in feeding days influence our biomass modeling through our use of TGC in equation (5.1) to predict mean weight development. Moreover, the mortality rate will increase by $\hat{\mu}_{it}$, meaning the mortality rate μ_{it} is revised according to

$$\mu_{it} = \bar{\mu}_{it} + \hat{\mu}_{it} \cdot \zeta_{it}, \tag{5.16}$$

where $\bar{\mu}_{it}$ is the base mortality rate when no lice treatment occurs. The change in mortality caused by lice treatments enter our biomass model through the survivability relationship provided in equation (5.2) and as a consequence our definition of A_t given in equation (5.3).

5.3 Weight Distribution

The harvest weight of salmon is of seminal importance, since both the price achieved per kilogram and EBIT per kilogram are functions of the salmon weight. At the time of harvesting the fish have their weight distributed according to a probability density function (PDF) f. We discretize the continuous weight domain into discrete classes W_α representing the weights in the interval $[\underline{W}_\alpha, \overline{W}_\alpha)$. However, as price is our main interest there is a small caveat. Namely, that salmon are priced according to its head-on-gutted (HOG) weight. We assume that the HOG weight is a constant fraction κ of the salmon's sea weight at the time of harvest. As such we correct our weight intervals to be $\left[\frac{\underline{W}_\alpha}{\kappa}, \frac{\overline{W}_\alpha}{\kappa}\right)$. Letting F denote the cumulative distribution function (CDF) of f, we can calculate the fraction of total biomass Z^α that falls into weight class W_α through

$$\mathrm{Z}^\alpha = \frac{\int_{\underline{W}_\alpha/\kappa}^{\overline{W}_\alpha/\kappa} x f(x) dx}{\int_{-\infty}^{\infty} x f(x) dx} = \left(F\left(\frac{\overline{W}_\alpha}{\kappa}\right) - F\left(\frac{\underline{W}_\alpha}{\kappa}\right)\right) \frac{\mathrm{E}\left[X \,\middle|\, \frac{\underline{W}_\alpha}{\kappa} \leq X \leq \frac{\overline{W}_\alpha}{\kappa}\right]}{\mathrm{E}[X]}. \tag{5.17}$$

As smaller fish are usually filtered out before deployment, the weight distribution of a cohort at deployment follows a log-normal distribution, which is maintained through the cohort's rearing cycle. In the following we assume that the weight of a cohort is distributed according to a log-normal distribution parametrized according to

$$X \sim \text{Lognormal}(\mu, \sigma) = e^{\mu + \sigma S}, \tag{5.18}$$

where S is distributed according to a standard normal distribution

$$S \sim \mathcal{N}(0,1) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}s^2}. \tag{5.19}$$

Letting w denote the average weight of a cohort with a coefficient of variation CV, we have the relationships

$$\mu = \ln\left(\frac{w^2}{\sqrt{w^2 + (\mathrm{CV} \cdot w)^2}}\right), \tag{5.20}$$

and

$$\sigma^2 = \ln\left(1 + \mathrm{CV}^2\right). \tag{5.21}$$

Letting Φ denote the CDF of the standard normal distribution we obtain the following expression by applying equation (5.17) to a log-normal distribution using the above parametrization

$$\mathrm{Z}^{\alpha} = \Phi\left(\frac{\ln(\frac{\overline{W}_{\alpha}}{\kappa}) - \mu - \sigma^2}{\sigma}\right) - \Phi\left(\frac{\ln(\frac{\underline{W}_{\alpha}}{\kappa}) - \mu - \sigma^2}{\sigma}\right). \tag{5.22}$$

We add that this approach is flexible, in that by adjusting κ one can obtain other weight class measures such as sea weight and head-off-gutted weight. Moreover, similar results can also be obtained for the normal distribution under the same assumptions, and other distributions could be accommodated if the parameters that define the distribution are available.

5.4 Relative Prices

As explained earlier, salmon prices depend on the weight of the salmon. It is difficult to predict the future salmon price, but as Asche and Guttormsen (2001) and Forsberg and Guttormsen (2006) show there are persistent seasonal patterns in relative prices. We compute the relative prices with the weight class 4 – 5 kg as the reference class, based on data from the NASDAQ Salmon Index between 2013 and 2020. We find that patterns in relative price still exist as shown in Figure 5.2. The choice of 4 – 5 kg as the reference class is motivated by the fact that it carries the largest volumes. In Figure 5.3 we show monthly relative prices broken down per year. The data for November

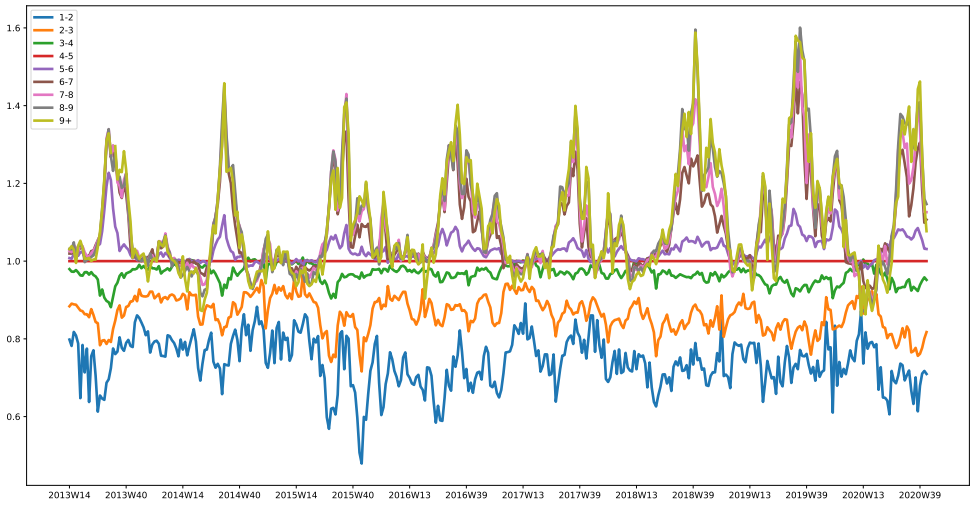

Figure 5.2: Weekly relative prices computed based on the NASDAQ Salmon Index between 2013 and 2020.

and December 2020 are lacking as the NASDAQ Salmon Index removed historic data per weight class from their website. The pattern of heavier weight classes' relative price peaking around August and September as Asche and Guttormsen (2001) found is still present. However, the pattern where the relative price of smaller salmon peak during summer and bottoms out during winter is no longer present, at least not consistently. For completeness, we mention that we also looked for seasonal patterns in absolute prices as found in Forsberg and Guttormsen (2006), however, we could not find any clear and persistent patterns across the years 2013 – 2020. Thus, we decide to only consider relative prices in this thesis as they historically have been rather predictable.

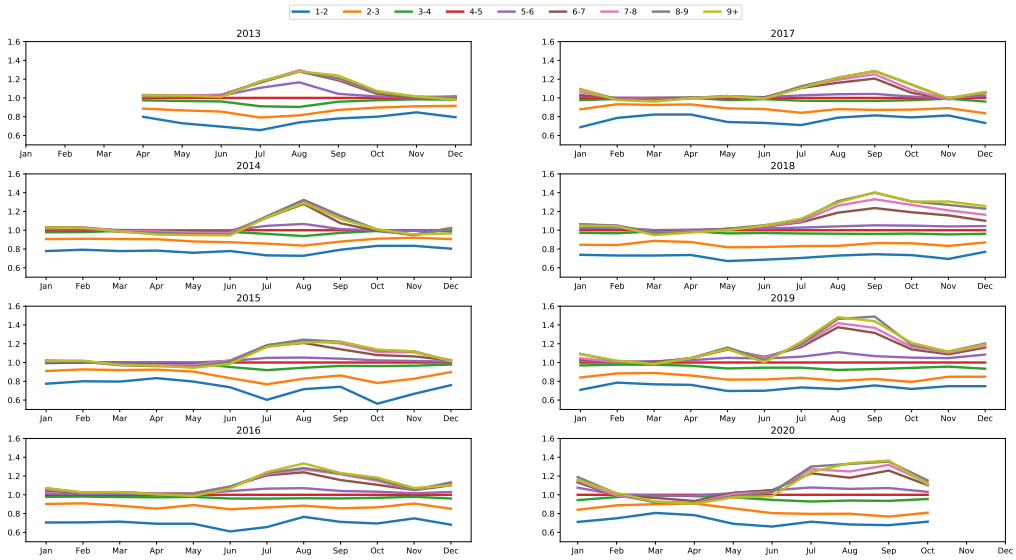

Figure 5.3: Monthly relative prices computed based on the NASDAQ Salmon Index between 2013 and 2020.

Chapter 6

Mathematical Model

We formulate the salmon farming company's decision problem as a multi-stage stochastic mixed integer program. In formulating this model we try to give a realistic depiction of the real world's requirements and constraints, while trying to provide a simple and clear model mathematically. We start off by giving an overview of the notation used in this chapter. Next, we describe the stages of our stochastic program. Thereafter, we introduce the program's objective function and constraints.

6.1 Notation

We now introduce the notation we use to formulate our mathematical model.

6.1.1 Sets

Symbol	Definition
$\mathcal{T}$	The set of all time periods in the planning horizon.
$\mathcal{T}^+$	As $\mathcal{T}$, but includes an extra dummy time period at the end of the planning horizon.
$\mathcal{T}^{\mathcal{R}}$	The set of all release time periods in the planning horizon.
$\mathcal{T}_0^{\mathcal{R}}$	$\mathcal{T}^{\mathcal{R}} \cup \{0\}$ where 0 is an initial dummy time period to account for initial conditions.

Symbol	Definition
$\mathcal{G}$	The set of all smolt types.
$\mathcal{G}_t$	The set of all smolt types that can be deployed during time period t.
$\mathcal{F}$	The set of all available initial smolt weights.
$\mathcal{F}_t$	The set of all available initial smolt weights that can be deployed during time period t.
$\mathcal{I}$	The set of all locations.
$\mathcal{T}_{fgit}^{\mathrm{G}^s}$	The set of all growth time periods in the planning horizon for a cohort with initial mean weight f and smolt type g deployed at location i during time period t in scenario s.
$\mathcal{T}_{fgit}^{\mathrm{H}^s}$	The set of all harvest time periods in the planning horizon for a cohort with initial mean weight f and smolt type g deployed at location i during time period t in scenario s.
$\mathcal{T}_{fgit}^{\mathrm{H}+^s}$	As $\mathcal{T}_{fgit}^{\mathrm{H}^s}$, but includes an extra final dummy time period in case the harvestable period starts before and continues after the end of the planning horizon.
$\mathcal{T}_{fgit}^{\mathrm{D}^s}$	The set of all time periods in scenario s when a cohort with initial mean weight f and smolt type g deployed at location i that can be harvested at time t could have been deployed.
$\mathcal{T}_t^{\Lambda^-}$	The set of the trailing Λ time periods at time t, including 0 when applicable.
$\mathcal{T}_t^{\Delta_R^-}$	The set of all release time periods occouring at time t or during the trailing $\Delta - 1$ time periods, including the dummy time period 0 when applicable.
$\mathcal{T}_t^{\Gamma^-}$	The set including the time period t and and the $\Gamma - 1$ trailing time periods.
$\mathcal{T}_i^{\Gamma^{\mathrm{INIT}}}$	The set including the time periods wherin there must be deployed biomass at location i during the start of the planning horizon to not be in conflict with the activity regulations.

Symbol	Definition
$\mathcal{T}^{\mathrm{E}^s}$	The set of time periods from the time period following the end of the planning horizon until the last time period where a cohort enjoys its growth phase in scenario s.
$\mathcal{W}$	The set of all HOG harvest weight classes.
$\mathcal{P}$	A set containing sets of time periods where smolt supply is restricted during the time periods in the set.
$\mathcal{B}^{\mathrm{G}^s}_{it}$	Contains tuples (f, g, t') representing deployments that would still be in their growth phase during time period t at location i in scenario s if deployed.
$\mathcal{B}^{\mathrm{H}^s}_{it}$	Contains tuples (f, g, t') representing deployments that would be in their harvest phase during time period t at location i in scenario s if deployed.
$\mathcal{S}$	The set of scenarios.
$\mathcal{N}$	The set of nodes in the scenario tree.
$\mathcal{S}_n$	The set of scenarios passing through node n in the scenario tree.
$\mathcal{T}_n$	The set of time periods associated with node n in the scenario tree.
$\mathcal{T}^{\mathcal{R}}_n$	The set of release time periods associated with node n in the scenario tree.

6.1.2 Parameters

Symbol	Definition
π^s	The probability of scenario s occurring.
P^{α}_{t}	The price of weight class α during time period t.
$\mathrm{Z}^{\alpha^s}_{t'fgit}$	The fraction of biomass from a deployment with initial mean weight f and smolt type g deployed at location i during time period t' that falls into HOG weight class α during time period t.

Symbol	Definition
N_f	Number of fish in thousands per tonne of smolt deployed from the release weight class f.
$\tilde{\mathrm{L}}_i^{\mathrm{DEP}}$	Lower bound on the number of fish that can be deployed simultaneously as part of a deployment of a specific smolt type at location i.
$\hat{\mathrm{L}}_i^{\mathrm{DEP}}$	Lower bound on the number of fish that can be deployed simultaneously at location i.
$\tilde{\mathrm{U}}_i^{\mathrm{DEP}}$	Upper bound on the number of fish that can be deployed simultaneously as part of a deployment of a specific smolt type at location i.
$\hat{\mathrm{U}}_i^{\mathrm{DEP}}$	Upper bound on the number of fish that can be deployed simultaneously at location i.
$\mathrm{L}_i^{\mathrm{H}}$	Lower bound on the amount of biomass that can be harvested simultaneously at location i.
$\mathrm{U}_i^{\mathrm{H}}$	Upper bound on the amount of biomass that can be harvested simultaneously at location i.
$\mathrm{U}_t^{\mathrm{COMP}}$	Upper bound on the amount of biomass that can be harvested at time t on a company wide level.
Q_i	Lower bound on the amount of biomass that must be employed at location i if biomass is employed.
$\mathrm{A}^s_{t'fgit}$	The ratio between the employed biomass during time period t and the initially deployed biomass during time period t' of a deployment with initial mean weight f and smolt type g deployed at location i during time period t'.
$\mathrm{R}^s_{t'fgit}$	The ratio between the employed biomass during time period t and $t-1$ of a deployment with initial mean weight f and smolt type g deployed at location i during time period t'.
MAB_i	Location Maximum Allowed Biomass (MAB) for location i.
$\mathrm{MAB}^{\mathrm{COMP}}$	Company wide Maximum Allowed Biomass (MAB).

Symbol	Definition
q_{0fgi0}	Initial amount of biomass originating from a smolt deployment with initial weight class f and smolt type g at location i at the start of the planning horizon.
y_{fgi0}	Initial amount of biomass originating from a smolt deployment with initial weight class f and smolt type g at location i at the start of the planning horizon.
β_{i0}	1 if and only if biomass is employed at location i at the start of the planning horizon.
Γ	The maximum amount of time periods without biomass presence before a license is vacated.
Δ	The maximum length of the rearing cycle.
Λ	The minimum amount of time periods a location must be fallowed before deploying new biomass after the end of a rearing cycle.

6.1.3 Decision Variables

Symbol	Definition
$\tilde{\delta}^s_{fgit}$	1 if and only if a cohort with initial mean weight f of type g is deployed at location i during time period t in scenario s, otherwise 0.
$\tilde{\delta}'^n_{fgit}$	Used to enforce non-anticipativity variable for $\tilde{\delta}^s_{fgit}$ in scenario tree node n if scenario s passes through n.
$\hat{\delta}^s_{it}$	1 if and only if a cohort is deployed at location i during time period t in scenario s, otherwise 0.
$\hat{\delta}'^n_{it}$	Used to enforce non-anticipativity variable for $\hat{\delta}^s_{it}$ in scenario tree node n if scenario s passes through n.
y^s_{fgit}	Continuous variable denoting the amount of biomass deployed through a cohort with initial weight f and type g at location i during time period t in scenario s.

Symbol	Definition
y'^{n}_{fgit}	Used to enforce non-anticipativity variable for y^{s}_{fgit} in scenario tree node n if scenario s passes through n.
ω^{s}_{it}	1 if and only if a harvesting operation takes place at location i during time period t in scenario s, otherwise 0.
$w^{s}_{t'fgit}$	Amount of tonnes harvested during time period t from a cohort with mean initial weight f, smolt type g, deployed at location i during time period t' in scenario s.
$w'^{n}_{t'fgit}$	Used to enforce non-anticipativity variable for $w^{s}_{t'fgit}$ in scenario tree node n if scenario s passes through n.
β^{s}_{it}	1 if and only if biomass is employed at location i during time period t in scenario s, otherwise 0.
β'^{n}_{it}	Used to enforce non-anticipativity variable for β'^{n}_{it} in scenario tree node n if scenario s passes through n.
$\tilde{q}^{s}_{t'fgit}$	The amount of biomass in tonnes at time t stemming from a cohort with initial mean weight f, smolt type g and deployed at location i during time period t' in scenario s.
$\hat{q}^{s}_{it}$	The amount of biomass in tonnes employed at location i during time period t in scenario s.

6.2 Stage Description

In this section we describe the stages of our multi-stage program. In general, the company plans her operations for a set of upcoming contiguous time periods at a time, e.g. for the upcoming year, which we denote a planning block. Before a planning block commences, the company must have decided all deployment variables, and when and where to fallow and harvest for all time periods contained in the planning block. The company makes these decisions assuming one of the scenarios $s \in \mathcal{S}$ will occur, but she does not know which one. Importantly, she does not need to decide the precise harvest volumes a head of time, she can rather wait for updated information, i.e. learn which scenario transpires, before she decides the harvest volumes for the time periods in the planning block. Assuming the company's planning horizon consists of one planning block spanning the entire horizon, we would

have two stages which decisions are made in.

In our case we wish to plan for a planning horizon containing multiple consecutive disjoint planning blocks, in the extreme case each time period could be its own planning block. Naturally, we get more than two stages, as the decisions made for each planning block should be based on the most updated information possible. Hence, the first-stage consists of deciding deployment variables and when and where to harvest for the first planning block, while the second-stage consists of deciding the harvest-volumes for the time periods of the first planning block knowing which scenario has transpired during the first planning block, and the third stage consists of deciding the deployment variables and when and where to harvest for the second planning block knowing the scenario that have transpired during the first planning block, and so on. As the second and third stage decisions are made having the same information available, these two stages are actually the same stage, and are thus merged. However, note that the very last stage cannot be merged with another stage, it therefore only consists of deciding the harvest volumes for the last planning block. Hence, a problem with n planning blocks has $n + 1$ stages.

Figure 6.1 illustrates a case with a planning horizon consisting of two time periods, and each time period is its own planning block for the sake of simplicity. Moreover, the figure depicts two possible scenarios per planning block, yielding a total of four scenarios. Each level of the tree represents a stage of our problem, and the variables inside the nodes are the variables that must be decided with the information available at that stage of the problem. The arrows indicate the information that becomes available between the stages, namely whether a lice treatment takes place or not.

The tree presented in Figure 6.1 can more generally be though of as a scenario tree. We let each node of the scenario tree be associated with the time periods for which deployment variables must be decided in the node, which we denote $\mathcal{T}_n$ for node n. Note that the leaves of the tree are thus not associated with any time periods.

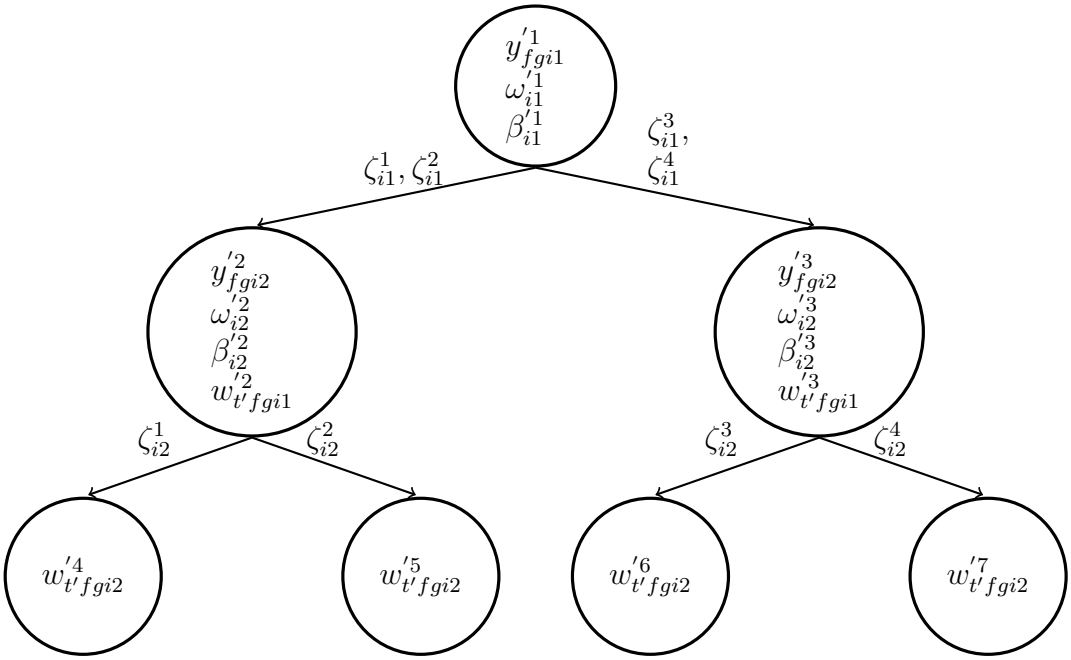

Figure 6.1: Depiction of the different stages of a problem consisting of two time periods and two scenarios per time period, where each time period is its own planning block.

6.3 Objective Function

We let our objective value be the expected value of total harvest volumes weighted by relative price. The product $\mathrm{Z}^{\alpha^s}_{t'fgit} w^s_{t'fgit}$ gives the harvest volumes contained in HOG weight class α for a harvest operation taking place during time period t, harvesting a cohort of smolt type g and initial mean weight f deployed during time period t' at location i. To get the total harvest volumes of such a cohort we sum over all time periods t contained in its harvest phase $\mathcal{T}^{\mathrm{H}^s}_{fgit'}$. Hence we arrive at the following objective

$$\max \sum_{s\in\mathcal{S}} \pi^s \left[\sum_{t'\in\mathcal{T}^{\mathcal{R}}0} \sum_{f\in\mathcal{F}_{t'}} \sum_{g\in\mathcal{G}_{t'}} \sum_{i\in\mathcal{I}} \sum_{t\in\mathcal{T}^{\mathrm{H}^s}_{fgit'}} \sum_{\alpha\in\mathcal{W}} \mathrm{P}^{\alpha}_{t} \mathrm{Z}^{\alpha^s}_{t'fgit} w^s_{t'fgit} \right]. \tag{6.1}$$

6.4 Constraints

As laid out in section 2.6 farming companies must abide by several practical and regulatory constraints. We now introduce the constraints we consider in this thesis. For the sake of brevity we discuss most of the constraints for a isolated single scenario, save the non-anticipativity constraints.

6.4.1 Smolt Deployments

Practical considerations require that a cohort deployed at location i contains at least $\tilde{\mathrm{L}}_i^{\mathrm{DEP}}$ fish, and no more than $\tilde{\mathrm{U}}_i^{\mathrm{DEP}}$ fish. The number of fish contained in a cohort is given by $\mathrm{N}_f \cdot y_{fgit}^s$, or in plain words the number of fish per tonne times the amount of tonnes deployed. We thus have the following constraint

$$\tilde{\mathrm{L}}_i^{\mathrm{DEP}} \tilde{\delta}_{fgit}^s \leq \mathrm{N}_f y_{fgit}^s \leq \tilde{\mathrm{U}}_i^{\mathrm{DEP}} \tilde{\delta}_{fgit}^s \quad s \in \mathcal{S}, t \in \mathcal{T}, f \in \mathcal{F}_t, g \in \mathcal{G}_t, i \in \mathcal{I}. \tag{6.2}$$

This constraint also enforces that $\tilde{\delta}_{fgit}^s = 1$ if and only if $y_{fgit}^s > 0$ i.e. when a corresponding deployment takes place, assuming $\tilde{\mathrm{L}}_i^{\mathrm{DEP}} > 0$.

In a similar manner, we require that if multiple cohorts are deployed at location i at the same time, they must in sum consist of at least $\hat{\mathrm{L}}_i^{\mathrm{DEP}}$ fish and at a maximum of $\hat{\mathrm{U}}_i^{\mathrm{DEP}}$ fish.

$$\hat{\mathrm{L}}_i^{\mathrm{DEP}} \hat{\delta}_{it}^s \leq \sum_{f \in \mathcal{F}_t} \sum_{g \in \mathcal{G}_t} \mathrm{N}_f y_{fgit}^s \leq \hat{\mathrm{U}}_i^{\mathrm{DEP}} \hat{\delta}_{it}^s \quad s \in \mathcal{S}, i \in \mathcal{I}, t \in \mathcal{T} \tag{6.3}$$

The reasoning for restricting smolt deployments at both a cohort and location level is of practical nature. The lower bound on a cohort level allows us to avoid very small smolt orders for a specific smolt type, and the location level lower bound ensures that there is enough smolt deployed to justify the use of operational resources required to operate the location. In practice, we only need an upper bound at a location level to ensure that the amount of smolt does not exceed operational limitations. The upper bound at the cohort level is added to ensure that $y_{fgit}^s = 0$ when $\tilde{\delta}_{fgit}^s = 0$.

The company has her smolt supply limited during certain set of time periods p, typically specific parts of the year. Meaning that the sum of deployed smolt across all cohorts deployed during the time periods contained in p must

be less than the supply S_p. Hence the following must hold

$$\sum_{i\in\mathcal{I}}\sum_{t\in p}\sum_{g\in\mathcal{G}_t}\sum_{f\in\mathcal{F}_t}\mathrm{N}_f y^s_{fgit} \leq \mathrm{S}_p \quad s\in\mathcal{S}, p\in\mathcal{P}. \tag{6.4}$$

To tighten the model we add the following constraint that forces the location deployment variable $\hat{\delta}^s_{it}$ to one, if at least one of the cohort deployment variables $\tilde{\delta}^s_{fgit}$ is equal to one.

$$\sum_{f\in\mathcal{F}_t}\sum_{g\in\mathcal{G}_t}\tilde{\delta}^s_{fgit} \leq |\mathcal{F}_t||\mathcal{G}_t|\hat{\delta}^s_{it} \quad s\in\mathcal{S}, i\in\mathcal{I}, t\in\mathcal{T}^{\mathcal{R}} \tag{6.5}$$

The product $|\mathcal{F}_t|\cdot|\mathcal{G}_t|$ represents the maximal amount of cohorts that can be deployed at a location during time period t, and is chosen as the right hand side coefficient to make the model as tight as possible.

When we deploy a cohort it must be at sea for at least the duration of the cohort's growth period. As such we must require β^s_{it} to be 1 during the cohort's growth period. For each possible deployment represented through $\tilde{\delta}^s_{fgit}$ we have

$$\tilde{\delta}^s_{fgit} \leq \beta^s_{i\tau} \quad s\in\mathcal{S}, t\in\mathcal{T}^{\mathcal{R}}, f\in\mathcal{F}_t, g\in\mathcal{G}_t, i\in\mathcal{I}, \tau\in\mathcal{T}^{\mathrm{G}^s}_{fgit}. \tag{6.6}$$

Together with constraint (6.15) this forces β^s_{it} to 1 when a location i contains biomass at time t. In fact constraint (6.15) is sufficient in itself as it enforces the equivalence that $\beta^s_{it}=1$ if and only if biomass is employed at location i during time period t. However, we find that the addition of constraint (6.6) makes the model tighter.

6.4.2 Harvesting

Similarly to deployment volumes, harvesting volumes are limited at a location level due to practical considerations and capacity constraints. To obtain the harvest volumes during time period t at location i, we must sum over all possible cohort's that would be in their harvest phase if deployed. Hence we consider all possible combinations of the initial mean weight f and smolt type g, and for each such combination we exploit that the set $\mathcal{T}^{\mathrm{D}}_{fgit}$ contains all possible release time periods that would yield a cohort in its harvest phase

at time t respecting the combination of f and g. We then sum over all harvest volume variables $w^s_{t'fgit}$ corresponding to a cohort which would be in its growth phase at time t

$$\mathrm{L}^{\mathrm{H}}_i \omega^s_{it} \leq \sum_{f\in\mathcal{F}} \sum_{g\in\mathcal{G}} \sum_{t'\in\mathcal{T}^{\mathrm{D}}_{fgit}} w^s_{t'fgit} \leq \mathrm{U}^{\mathrm{H}}_i \omega^s_{it} \quad s\in\mathcal{S}, i\in\mathcal{I}, t\in\mathcal{T}. \tag{6.7}$$

This constraint also ensures that $\omega^s_{it} = 1$ if and only if harvesting takes place, given $\mathrm{L}^{\mathrm{H}}_i > 0$.

The company must also abide by a company wide harvest limit during a time period $\mathrm{U}^{\mathrm{COMP}}_t$. The company wide harvest volumes are obtained through simply summing up the harvest volumes at each individual location.

$$\sum_{f\in\mathcal{F}} \sum_{g\in\mathcal{G}} \sum_{i\in\mathcal{I}} \sum_{t'\in\mathcal{T}^{\mathrm{D}}_{fgit}} w^s_{t'fgit} \leq \mathrm{U}^{\mathrm{COMP}}_t \quad s\in\mathcal{S}, t\in\mathcal{T} \tag{6.8}$$

We note that the harvest volume constraints are good targets for relaxation, as they are in reality soft constraints, e.g. going one kilogram above or below the harvest volume constraints is insignificant from a real world perspective.

6.4.3 Fallowing

The company is required to fallow a location after the end of a rearing cycle. Denoting the required fallowing length in time periods as Λ, we must ensure that no deployment can take place at a location i unless there were no biomass present at this location during the trailing Λ time periods. If a deployment takes place at time t we have that $\hat{\delta}^s_{it} = 1$, and to uphold the fallowing requirement we must have $\beta^s_{i\tau} = 0$ for all $\tau \in \{t-\Lambda, t-\Lambda+1, \ldots, t-1\} = \mathcal{T}^{\Lambda^-}_t$, yielding

$$\Lambda\hat{\delta}^s_{it} + \sum_{\tau\in\mathcal{T}^{\Lambda^-}_t} \beta^s_{i\tau} \leq \Lambda \quad s\in\mathcal{S}, i\in\mathcal{I}, t\in\mathcal{T}^{\mathcal{R}} \setminus \{1,..,\Lambda\}. \tag{6.9}$$

If $\hat{\delta}^s_{it} = 1$ in the restriction above, it must be the case that β^s_{it} is 0 for the trailing Λ time periods. Thus, there must have been zero biomass employed at the location during the trailing Λ time periods, as we shall enforce that $\beta^s_{it} = 0$ if and only if there is no biomass present at location i during time period t.

6.4.4 Activity Requirements

To avoid vacating a license, the company must ensure that she has employed at least some biomass at each location during a period consisting of Γ consecutive time periods. As $\beta^s_{it} = 1$ if and only if biomass is employed at location i during time period t, it suffices that $\beta^s_{i\tau} = 1$ for at least one $\tau \in \{\tau - \Gamma + 1, \tau - \Gamma + 2, \ldots, \tau\} = \mathcal{T}_\tau^{\Gamma^-}$, for all contiguous sequences of Γ time periods.

$$\sum_{\tau \in \mathcal{T}_t^{\Gamma^-}} \beta^s_{i\tau} \geq 1 \quad s \in \mathcal{S}, i \in \mathcal{I}, t \in \mathcal{T} \setminus \{1, ..., \Gamma - 1\} \tag{6.10}$$

6.4.5 Biomass Development and Regulations

We keep track of biomass development at both a cohort, location, and company level. The basic building block is the cohort level. We then aggregate the cohort level into a location level biomass, and thereafter the locations' biomass are aggregated into the company's total employed biomass.

Cohort Level

We must keep track of a cohort's total biomass as it develops over time, as growth and mortality takes place. Since the biomass development is discretized with respect to time, we assume that growth and mortality occur at the start of the time period following the time period where they were accrued. To illustrate this, a cohort deployed during time period t' accrues biomass development during time period t', but this development is not accounted for before the start of time period $t' + 1$. Moreover, we assume that all deployments take place during the start of the time period it is performed.

The simplest case of biomass development we need to account for is when no harvesting can take place. Due the simplicity of this case, we separate it from the harvest case to ease the computational strain. The parameter $\mathrm{A}^s_{t'fgit}$ denotes the ratio between a cohort's total biomass at time t and its initial biomass at the time of deployment t', we introduced this parameter in section 5.1 with a more modest amount of subscripts. When no harvesting takes place we can obtain the total biomass of a cohort at time t, $\tilde{q}^s_{t'fgit}$, through multiplying the initially deployed biomass $y^s_{fgit'}$ with the ratio $\mathrm{A}^s_{t'fgit}$.

This approach is valid during the entire growth phase and the first time period of the harvest phase. The latter is included to account for the biomass development accrued during the last time period of the growth phase.

$$\tilde{q}^s_{t'fgit} = \mathrm{A}^s_{t'fgit} y^s_{fgit'} \quad \begin{array}{l} s \in \mathcal{S}, t' \in \mathcal{T}_0^{\mathcal{R}}, f \in \mathcal{F}_{t'}, g \in \mathcal{G}_{t'}, i \in \mathcal{I}, \\ t \in \mathcal{T}^{\mathrm{G}^s}_{fgit'} \cup \{\min \mathcal{T}^{\mathrm{H}^s}_{fgit'}\} \end{array} \tag{6.11}$$

During the harvest phase things are slightly more complicated as we have to account for the harvest volumes $w^s_{t'fgit}$. $\mathrm{R}^s_{t'fgit}$ is a parameter giving the ratio between a cohort's total biomass at time t and $t-1$ if no harvesting takes place. If a harvest operation takes place during time t, we assume that it takes place at the start of the time period. As a consequence, the harvested fish will not have accrued any biomass growth or mortality. Furthermore, we assume that the company uses a batch harvesting scheme as described in Forsberg (1999), meaning that it will not filter fish within a cohort by weight when harvesting. Under these assumptions the total biomass at time $t+1$ is simply the biomass development ratio $\mathrm{R}^s_{t'fgi(t+1)}$ times the biomass at the start of time period t with the volumes that are harvested immediately at the start of the time period subtracted.

$$\tilde{q}^s_{t'fgit} = \mathrm{R}^s_{t'fgit}(\tilde{q}^s_{t'fgi(t-1)} - w^s_{t'fgi(t-1)}) \quad \begin{array}{l} s \in \mathcal{S}, t' \in \mathcal{T}_0^{\mathcal{R}}, f \in \mathcal{F}_{t'}, g \in \mathcal{G}_{t'}, i \in \mathcal{I}, \\ t \in \mathcal{T}^{\mathrm{H}+^s}_{fgit'} \setminus \{\min \mathcal{T}^{\mathrm{H}^s}_{fgit'}\} \end{array} \tag{6.12}$$

Caution must be shown during the final time period $|\mathcal{T}|$ of the planning horizon. If the cohort's growth phase extends beyond the final time period, a final dummy time period $1+|\mathcal{T}|$ is included in the set $\mathcal{T}^{\mathrm{H}+^s}_{fgit'}$ to allow for harvesting during time period $|\mathcal{T}|$. The reader might note that we have still not allowed harvesting to take place during a cohort's last harvest time period, this issue is accounted for in constraint (6.13).

Next, we require that all deployed fish are eventually harvested. If this is not explicitly enforced the company might find it advantageous to not harvest the entire deployed cohort due to the harvest volume constraints. We note that without the harvest constraints, the company would always harvest every deployed fish as a consequence of the objective maximizing for harvest volumes. We enforce that the harvest volume during the last harvest time

period must be equal to whatever biomass that is left of the cohort.

$$\tilde{q}^{s}_{t'fgi\left[\max \mathcal{T}^{\text{H}+s}_{fgit'}\right]} - w^{s}_{t'fgi\left[\max \mathcal{T}^{\text{H}+s}_{fgit'}\right]} = 0 \quad s \in \mathcal{S}, t' \in \mathcal{T}^{\mathcal{R}}_{0}, f \in \mathcal{F}_{t'}, g \in \mathcal{G}_{t'}, i \in \mathcal{I} \tag{6.13}$$

Again, extra care must be shown at the end of the planning horizon. If a cohort's harvest period extends beyond $|\mathcal{T}|$, we avoid harvesting the entire cohort at time $|\mathcal{T}|$ through the extra final dummy time period $1 + |\mathcal{T}|$ in $\mathcal{T}^{\text{H}+s}_{fgit'}$. The harvest volumes stemming from the final dummy time period are of technical nature and are not part of the objective function. Thus, in a practical implementation the above constraint can be omitted for cohorts with growth periods extending beyond the planning horizon.

Location Level

We introduce $\hat{q}^{s}_{it}$ as a book keeping variable, that keeps track of the amount of biomass employed at location i during time period t. To obtain the total biomass at location i during time period t we must account for the biomass contributed by all deployments that could have been deployed during time t and the trailing $\Delta - 1$ time periods. This time window is sufficient as Δ is the maximum amount of time periods in a rearing cycle. However, we only need to consider the time periods in this time window that allow deployments $\mathcal{T}^{\Delta^{-}_{R}}_{t}$. Hence for each release time step $t' \in \mathcal{T}^{\Delta^{-}_{R}}_{t}$ we sum up the biomass at time t for each cohort that could have been deployed at time t' at location i.

$$\hat{q}^{s}_{it} = \sum_{t' \in \mathcal{T}^{\Delta^{-}_{R}}_{t}} \sum_{f \in \mathcal{F}_{t'}} \sum_{g \in \mathcal{G}_{t'}} \tilde{q}^{s}_{t'fgit} \quad s \in \mathcal{S}, i \in \mathcal{I}, t \in \mathcal{T} \tag{6.14}$$

The biomass at a location is bounded by MAB_i due to regulations. Moreover, we enforce the equivalence that $\beta^{s}_{it} = 1$ if and only if biomass is employed at location i at time t. To achieve this we introduce a lower bound on biomass Q_i when $\beta^{s}_{it} = 1$.

$$\text{Q}_i \beta^{s}_{it} \leq \hat{q}^{s}_{it} \leq \text{MAB}_i \beta^{s}_{it} \quad s \in \mathcal{S}, i \in \mathcal{I}, t \in \mathcal{T} \tag{6.15}$$

Company Level

Lastly, the company is subject to a company wide MAB. Thus, the sum of biomass across all locations must be less than or equal to MAB^{COMP} during

all time periods.

$$\sum_{i \in \mathcal{I}} \hat{q}^s_{it} \leq \text{MAB}^{\text{COMP}} \quad s \in \mathcal{S}, t \in \mathcal{T} \tag{6.16}$$

6.4.6 Initial Conditions

We must account for the fact that there is an operational history preceding our planning horizon. This means there is already biomass present at the start of the planning horizon. We account for this through the dummy time period 0, and the parameters y^s_{fgi0}, $\tilde{q}^s_{0fgi0}$, and β^s_{i0}. These deployments have their first biomass adjustment at the start of the second time period, meaning $\text{A}_{0fgi1} = 1$, and for simplicity that all the cohort's are in their growth phase during the dummy time period. Moreover, the set of harvest periods $\mathcal{T}^{\text{H}}_{fgit'}$ are adjusted to make sure that the dummy cohorts do not have a longer rearing cycle than Δ time periods, based on the cohort's time at sea preceding the planning horizon. We add that there is no need to keep the different initial mean weights separated for these dummy deployments. All cohorts can be assigned to the same initial mean weight class f. The A and R parameters for these cohorts are based on the weight and time spent at sea at the start of the planning horizon.

Lastly, we introduce a set of initial activity constraints which serve the role of constraint (6.10) at the start of the planning horizon, i.e. for $t \leq \Gamma$. The initial condition set $\mathcal{T}_i^{\Gamma^{\text{INIT}}}$ contains the time periods where the company must employ biomass during at least one time period. The set contains Γ contiguous time periods starting from the time period where location i was last fallowed, however, we discard the time periods preceding the planning horizon.

$$\sum_{t \in \mathcal{T}_i^{\Gamma^{\text{INIT}}}} \beta^s_{i\tau} \geq 1 \quad s \in \mathcal{S}, i \in \mathcal{I} \tag{6.17}$$

In a similar fashion we also introduce a set of initial fallowing constraints

$$\beta^s_{it} \leq 0 \quad s \in \mathcal{S} i \in \mathcal{I}, t \in \mathcal{T}_i^{\Lambda^{\text{INIT}}}. \tag{6.18}$$

6.4.7 End of Horizon Constraints

Lastly, the company does in reality plan to keep operations running on an infinite time horizon, and as such the model cannot completely disregard the

future outside of the planning horizon. We therefore require that there is a certain amount of biomass deployed at the end of the planning horizon. Specifically, we demand that the company's total biomass at sea during the final time period of the planning horizon $|\mathcal{T}|$ is greater than the total initial biomass

$$\sum_{i\in\mathcal{I}} \hat{q}^s_{i|\mathcal{T}|} \geq \sum_{f\in\mathcal{F}}\sum_{g\in\mathcal{G}}\sum_{i\in\mathcal{I}} y^s_{fgi0} \quad s\in\mathcal{S}. \tag{6.19}$$

The company must also be able to comply with the biomass regulations after the end of the planning horizon. We let $\mathcal{B}^{\mathrm{G}^s}_{it}$ contain the set of tuples (f,g,t') representing deployments of cohorts that would be in their growth phase at location i during time period t if deployed. $\mathcal{T}^{\mathrm{E}^s}$ contains the time periods from the end of the planning horizon up to and including the last possible growth period of a potentially deployed cohort. As these cohorts are in their growth phase at time t, their biomass at time t can be calculated as $\mathrm{A}^s_{t'fgit}y^s_{fgit'}$. We then arrive at

$$\sum_{(f,g,t')\in\mathcal{B}^{\mathrm{G}^s}_{it}} \mathrm{A}^s_{t'fgit}y^s_{fgit'} \leq \mathrm{MAB}_i \quad s\in\mathcal{S}, i\in\mathcal{I}, t\in\mathcal{T}^{\mathrm{E}^s} \tag{6.20}$$

$$\sum_{i\in\mathcal{I}}\sum_{(f,g,t')\in\mathcal{B}^{\mathrm{G}^s}_{it}} \mathrm{A}^s_{t'fgit}y^s_{fgit'} \leq \mathrm{MAB}^{\mathrm{COMP}} \quad s\in\mathcal{S}, t\in\mathcal{T}^{\mathrm{E}^s}. \tag{6.21}$$

6.4.8 Non-Anticipativity Constraints

We now introduce the non-anticipativity constraints. For each node in the scenario tree we force the decision variables corresponding to the same decision, e.g. ω^s_{it}, to the same value for all scenarios passing through the node. We represent the common value by a decision variable marked with $'$, e.g. ω'^n_{it} for the binary harvest variables in node n. For each node n we associate a set of time periods $\mathcal{T}_n$, for which we must decide deployment volumes and whether to fallow and harvest or not up front. Denoting the set of scenarios passing through node n as $\mathcal{S}_n$ we have

$$(\tilde{\delta}^s_{fgit}, y^s_{fgit}) = (\tilde{\delta}'^n_{fgit}, y'^n_{fgit}) \quad n\in\mathcal{N}, s\in\mathcal{S}_n, t\in\mathcal{T}^{\mathcal{R}}_n, f\in\mathcal{F}_t, g\in\mathcal{G}_t, i\in\mathcal{I}, \tag{6.22}$$

$$\hat{\delta}^s_{it} = \hat{\delta}'^n_{it} \quad n\in\mathcal{N}, s\in\mathcal{S}_n, t\in\mathcal{T}^{\mathcal{R}}_n, f\in\mathcal{F}_t, g\in\mathcal{G}_t, i\in\mathcal{I}, \tag{6.23}$$

and

$$(\omega^s_{it}, \beta^s_{it}) = (\omega'^n_{it}, \beta'^n_{it}) \quad n\in\mathcal{N}, s\in\mathcal{S}_n, i\in\mathcal{I}, t\in\mathcal{T}_n. \tag{6.24}$$

Adding non-anticipativity constraints on $\tilde{\delta}^s_{fgit}$ and $\hat{\delta}^s_{it}$ is not strictly necessary, as the non-anticipativity constraints put on y^s_{fgit} only leaves one possible assignment for the aforementioned variables.

The harvest volume variables $w^s_{t'fgit}$ have non-anticipativity constraints enforced in the stage following the stage where the corresponding binary harvest variables ω^s_{it} have their non-anticipativity constraints enforced. Hence, for a node n we require that the harvest volumes matches for the time periods associated with the node's predecessor $\pi(n)$

$$w^s_{t'fgit} = w'^n_{t'fgit} \quad n \in \mathcal{N}, s \in \mathcal{S}_n, t \in \mathcal{T}^+_{\pi(n)}, i \in \mathcal{I}, (f,g,t') \in \mathcal{B}^{\text{H}^s}_{it}. \tag{6.25}$$

We thus have that the root node does not impose any non-anticipativity constraints on $w^s_{t'fgit}$.

6.4.9 Non-Negativity and Binary Requirements

For completeness we state the non-negativity and binary requirements.

$$\omega^s_{it}, \beta^s_{it} \in \{0,1\} \quad s \in \mathcal{S}, i \in \mathcal{I}, t \in \mathcal{T} \tag{6.26}$$

$$\omega'^n_{it}, \beta'^n_{it} \in \{0,1\} \quad n \in \mathcal{N}, i \in \mathcal{I}, t \in \mathcal{T}_n \tag{6.27}$$

$$\hat{\delta}^s_{it} \in \{0,1\} \quad s \in \mathcal{S}, i \in \mathcal{I}, t \in \mathcal{T}^{\mathcal{R}} \tag{6.28}$$

$$\hat{\delta}'^n_{it} \in \{0,1\} \quad n \in \mathcal{N}, i \in \mathcal{I}, t \in \mathcal{T}^{\mathcal{R}}_n \tag{6.29}$$

$$\tilde{\delta}^s_{fgit} \in \{0,1\} \quad s \in \mathcal{S}, t \in \mathcal{T}^{\mathcal{R}}, f \in \mathcal{F}_t, g \in \mathcal{G}_t, i \in \mathcal{I} \tag{6.30}$$

$$\tilde{\delta}'^n_{fgit} \in \{0,1\} \quad n \in \mathcal{N}, t \in \mathcal{T}^{\mathcal{R}}_n, f \in \mathcal{F}_t, g \in \mathcal{G}_t, i \in \mathcal{I} \tag{6.31}$$

$$\hat{q}^s_{it} \geq 0 \quad s \in \mathcal{S}, i \in \mathcal{I}, t \in \mathcal{T} \tag{6.32}$$

$$w^s_{t'fgit} \geq 0 \quad s \in \mathcal{S}, t' \in \mathcal{T}^{\mathcal{R}}_0, f \in \mathcal{F}_{t'}, g \in \mathcal{G}_{t'}, i \in \mathcal{I}, t \in \mathcal{T}^{\text{H}^+{}^s}_{fgit'} \tag{6.33}$$

$$w'^n_{t'fgit} \geq 0 \quad n \in \mathcal{N}, t \in \mathcal{T}^+_{\pi(n)}, i \in \mathcal{I}, (f,g,t') \in \bigcup_{s \in \mathcal{S}_n} \mathcal{B}^{\text{H}^s}_{it} \tag{6.34}$$

$$y^s_{fgit} \geq 0 \quad s \in \mathcal{S}, t \in \mathcal{T}^{\mathcal{R}}, f \in \mathcal{F}_t, g \in \mathcal{G}_t, i \in \mathcal{I} \tag{6.35}$$

$$y'^n_{fgit} \geq 0 \quad n \in \mathcal{N}, t \in \mathcal{T}^{\mathcal{R}}_n, f \in \mathcal{F}_t, g \in \mathcal{G}_t, i \in \mathcal{I}, \tag{6.36}$$

$$\tilde{q}^s_{ft'git} \geq 0 \quad s \in \mathcal{S}, t' \in \mathcal{T}^{\mathcal{R}}_0, f \in \mathcal{F}_{t'}, g \in \mathcal{G}_{t'}, i \in \mathcal{I}, t \in \mathcal{T}^{\text{G}^s}_{fgit'} \cup \mathcal{T}^{\text{H}^+{}^s}_{fgit'} \tag{6.37}$$

Chapter 7

Solution and Evaluation Methodologies

The mathematical model introduced in chapter 6 is a multi-stage stochastic program. Ideally, we would like to solve this program with a time granularity of one month, and having each planning block representing a year and potentially hundreds of scenarios. In the case studies that we introduce in chapter 8 we use a time horizon of five years and have a total of 16 locations. Unfortunately, it turns out that solving these cases with six stages is computationally infeasible even for as few as 2 scenarios per planning block. In this chapter we introduce heuristics and methods to simplify the problem to make it computationally tractable. In short, we transform the multi-stage program into a simplified two-stage program, and evaluate the first-stage solution by applying a rolling horizon heuristic based simulation framework.

7.1 Two-Stage Simplification

As the multi-stage program introduced in chapter 6 is intractable even for a modest amount of scenarios, we decide to reduce it to a two-stage program. This is done by simply removing the non-anticipativity constraints following the first-stage. After solving the two-stage program we obtain a solution for the first-stage, which we denote a candidate solution. We also get a scenario specific solution to the second-stage.

The candidate solution, specifies where, when, what and how much to deploy

and when and where to fallow and harvest during the time periods associated with the first planning block, however it does not specify the harvest volumes during these time periods. From a practical standpoint, a farming company is mainly interested in the candidate solution, as the decisions contained therein have to be decided more or less immediately, while the company can wait and see how the real world evolves before deciding the second-stage variables. At the same time, the contribution of the second-stage variables to the objective value informs the company about the quality of the candidate solution. In fact, our objective only consists of second-stage variables. Thus, our objective measures how good second-stage decisions our candidate solution allows.

Although we have reduced the computational strain by turning the multi-stage problem into a two-stage problem, the amount of scenarios we can include is still limited. As such, relying solely on the objective value of the two-stage program to evaluate the quality of a candidate solution is insufficient, as we do not know how well the candidate solution will fare outside of the included scenarios.

Farming companies usually plan one year at a time, i.e. planning blocks spanning one year. Thus, the company would use the two-stage program at the start of the year to obtain the coming year's operational plan. When the year has passed it would run the program again with updated initial conditions to produce the operational plan for the next year, and so on. In general, the company is faced with a sequential decision problem. As such, when we evaluate a candidate solution we should perform multiple simulations of a sequential decision process. In the following we thus develop a rolling horizon heuristic inspired simulation framework to evaluate candidate solutions. We then apply this framework several times to the same candidate solution to get a more reliable estimate of its quality.

7.2 Evaluation Framework

To provide a single evaluation of a candidate solution, we need to decide the remaining decision variables of the planning horizon. Hence, we first need to decide the harvest volume variables of the planning block covered by the candidate solution. As the harvest volume variables depends on knowing the actual realization of the planning block, we need to simulate the outcome

of this planning block. Next, we must decide the decision variables of the next planning block and so on. This can be done by simply repeating the process we used for the first planning block. When all decision variables of the planning horizon are decided, we can compute an objective value. In the following we first provide a deeper background into the scheme presented above, before we describe the process of obtaining multiple evaluations of a candidate solution in a manner that makes it possible to compare different candidate solutions.

7.2.1 Rolling Horizon Heuristic

We now provide some formal background to our evaluation scheme, which is based on a Rolling Horizon Heuristic (RHH).The main idea behind an RHH is to divide the problem into several subproblems along the time axis and solve these problems in chronological order. Each subproblem consists of three disjoint contigiuous time period sets, the locked period $\mathcal{T}^{\mathcal{L}}$, the central period $\mathcal{T}^{\mathcal{C}}$, and the forecast period $\mathcal{T}^{\mathcal{F}}$, where $\mathcal{T}^{\mathcal{L}} \cup \mathcal{T}^{\mathcal{C}} \cup \mathcal{T}^{\mathcal{F}}$ is contigiuous. For each subproblem the locked period provides the initial conditions, while the central and forecast period contribute the constraints, decision variables, and objective function.

The decisions made during the central period are locked and carried over as initial conditions to the subsequent subproblem, since the central period of a subproblem becomes part of the locked period of the next subproblem. The forecast period is included in the subproblem to ensure that the decisions made during the central period account to at least some degree for their future consequences. Since these decisions are not carried over to the subsequent problems, it is common to relax constraints and integrality requirements during this period to ease the computational load further. Note that the forecast period does not necessarily make up the entire time line following the central period, but rather a contiguous subsequence that is sufficiently long to ensure we do not over optimize for the central period at the cost of the future.

Figure 7.1 gives a graphical overview of a simple application of an RHH scheme, where we having a planning horizon of five years, a central period of one year, and a forecast period of two years. Note that we include year 6 and 7 as part of the forecast period during the last iteration, even though

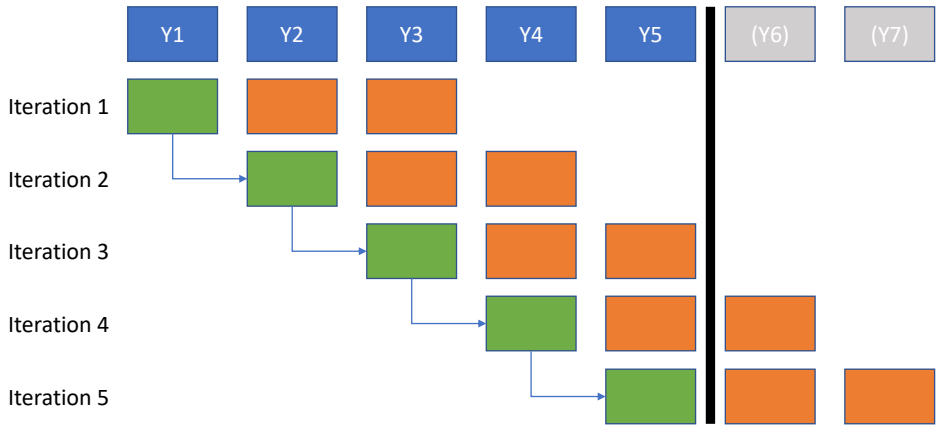

Figure 7.1: Execution of the rolling horizon heuristic over a five year planning horizon with a one year central period and a two year forecast period.

they are outside of the planning horizon.

7.2.2 Rolling Horizon Heuristic Combined With Simulation

We wish to apply an RHH scheme where each subproblem is made up of a two-stage stochastic program. In our case a simple RHH-scheme which breaks the planning horizon into different subproblems based solely on the time-axis is insufficient, as variables sharing the same time index are in different stages of our two-stage stochastic program. We decide to decompose the time-axis by planning blocks, such that each subproblem span the time periods contained in a planning block. However, we split each of these subproblems into two new subproblems, the first having a central period containing the first-stage variables of the planning block, and the second having a central period containing the second-stage variables of said planning block. As a consequence, our central periods and forecast periods are no longer disjoint with respect to time. Hence, when we move to the next central period, we do not necessarily move forward in time. In some contexts we simply refer to the second-stage variables of the planning block as harvest volume variables, this is a simplification as technically the second stage also contains the biomass decision variables $\tilde{q}$ and $\hat{q}$, but these are decided implicitly by the combination of deployment volume and harvest volume variables.

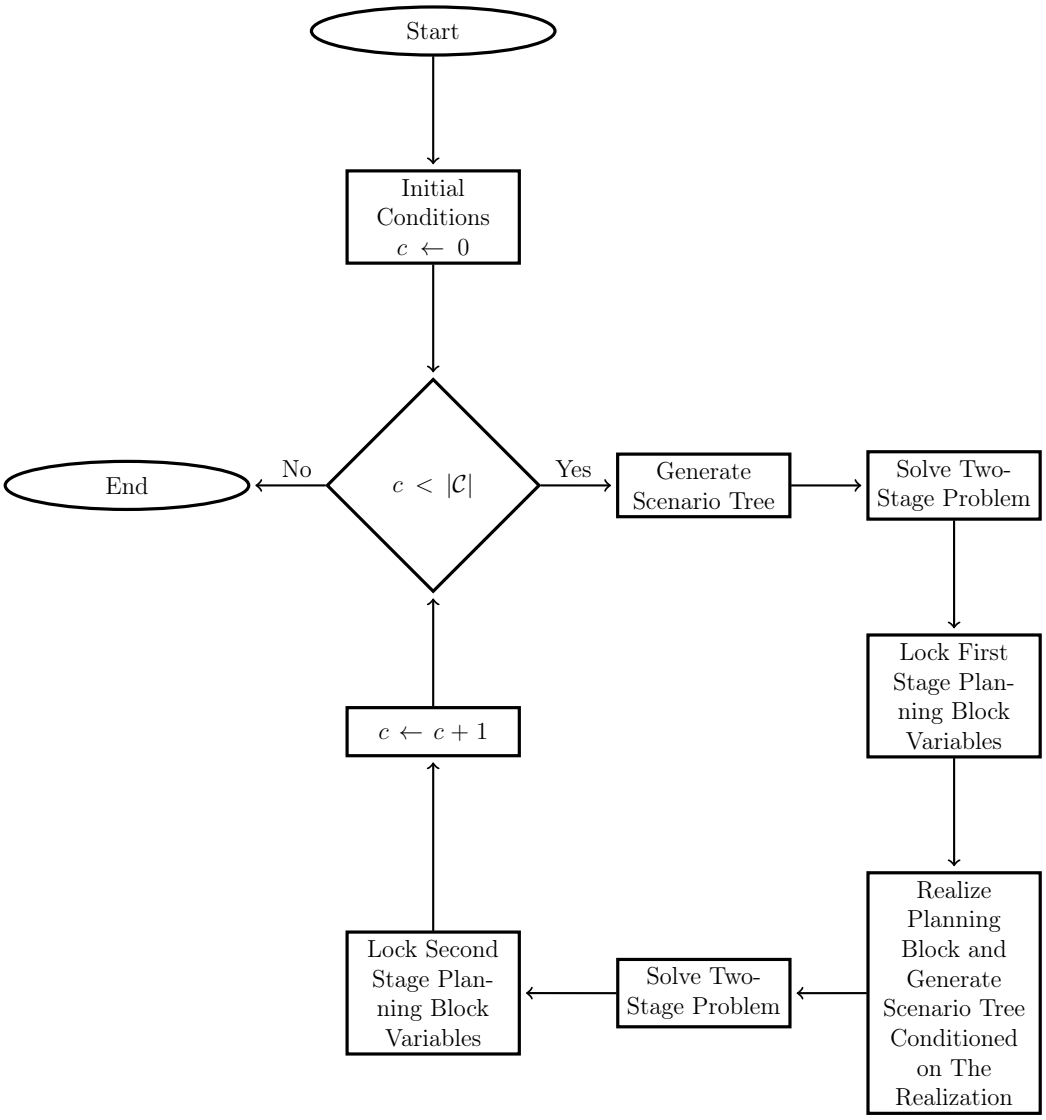

Figure 7.2: Schematic overview of the evaluation framework. $\mathcal{C}$ is the set of planning blocks.

For each sub-problem, we generate a two-stage stochastic program as described in section 7.1, by taking the model introduced in chapter 6 with a planning horizon spanning the time periods contained in the central and forecast period, and then relaxing the non-anticipativity constraints post the first-stage. This means our two-stage program also considers time periods after the forecast period, since we must account for the end of horizon constraints. We therefore deviate somewhat from the RHH philosophy, in that we solve complete replicas of the problem repeatedly, instead of just considering specific parts of the problem at a time.

Solving the subproblems containing first-stage variables is done through the same procedure as when we generate a candidate solution. In fact, generating a candidate solution can be viewed as solving the first subproblem of the RHH. The scenario tree used in this subproblem contains two levels, a root node for the central period, and leaf nodes for the forecast period. After obtaining the first-stage variables we lock them and move them to the locked period. Next, we simulate the outcome of the time periods spanned by the central period. To do this, we generate one specific outcome, i.e. one scenario, for the stochastic lice treatment variables for the time periods contained in the central period of the subproblem. We denote the scenario we generate as a realization of the planning block.

Next, we move onto the subproblem which solves for the harvest volume variables of the same time periods as the previous subproblem. However, we still wish to use stochastic programming to decide these variables, although we have a specific realization of the planning block. Hence, we generate a set of scenarios for the time periods contained in forecast period conditioned on our realization of the planning block time periods. Note that the central and forecast period are disjoint with respect to time in this subproblem. We choose to generate the same number of scenarios as we used when solving the first-stage variable subproblem. We thus get a two-stage scenario tree, which structure differs from the scenario tree of the previous subproblem in that branching has been delayed by one node, meaning the harvest volume variables of the central period have transitioned into the role of first-stage decision variables in this subproblem. Figure 7.3 illustrates this for a simple example with only two time periods each being a planning block, the tree on the left shows the scenario tree used during the first-stage subproblem, while

the tree on the right is used during the second-stage subproblem to solve for the harvest volume variables. We then solve the obtained subproblem, and lock the harvest volume variables.

The combination of having realized the planning block together with locking the harvest volume variables provide sufficient initial conditions for the next subproblem. The cohort level biomass variables, $\tilde{q}$, for the time period that follows the end of the planning block give the initial present biomass at the start of the next subproblem, with one important exception, we must discard the $\tilde{q}$ variables corresponding to cohorts deployed during the time period that immediately follows the planning block. Moreover, to be able to generate the A and R parameters, we must also store the mean weight of the cohorts at the start of the time period that immediately follows the planning block together with the time the cohort has spent at sea. Storing cohort's time at sea is also necessary to ensure that the cohort's rearing cycle length does not exceed Δ time periods.

We then move onto the next subproblem by moving to the next central period. In doing so, we also generate a new two-level scenario tree, where the root node is conditioned on the realization we generated for the previous planning block. This subproblem solves for the first-stage variables of the next planning block. We can now simply repeat this scheme, to produce a solution for the entire planning horizon. We summarize the main elements of this scheme in Figure 7.2.

As we now have generated a solution for the entire planning horizon, we need to compute the objective value of this solution. In our case this is done at the end of each iteration of the solution framework presented in Figure 7.2, i.e. after the harvest volume variables are locked. For each iteration we compute all terms in the objective provided in equation (6.1) containing a harvest volume variable $w^s_{t'fgit}$ with a time index t being part of the planning block, or more generally all terms connected to the planning block. After all iterations are finished, we sum up the objective values obtained during each iteration to get a total objective value for the planning horizon.

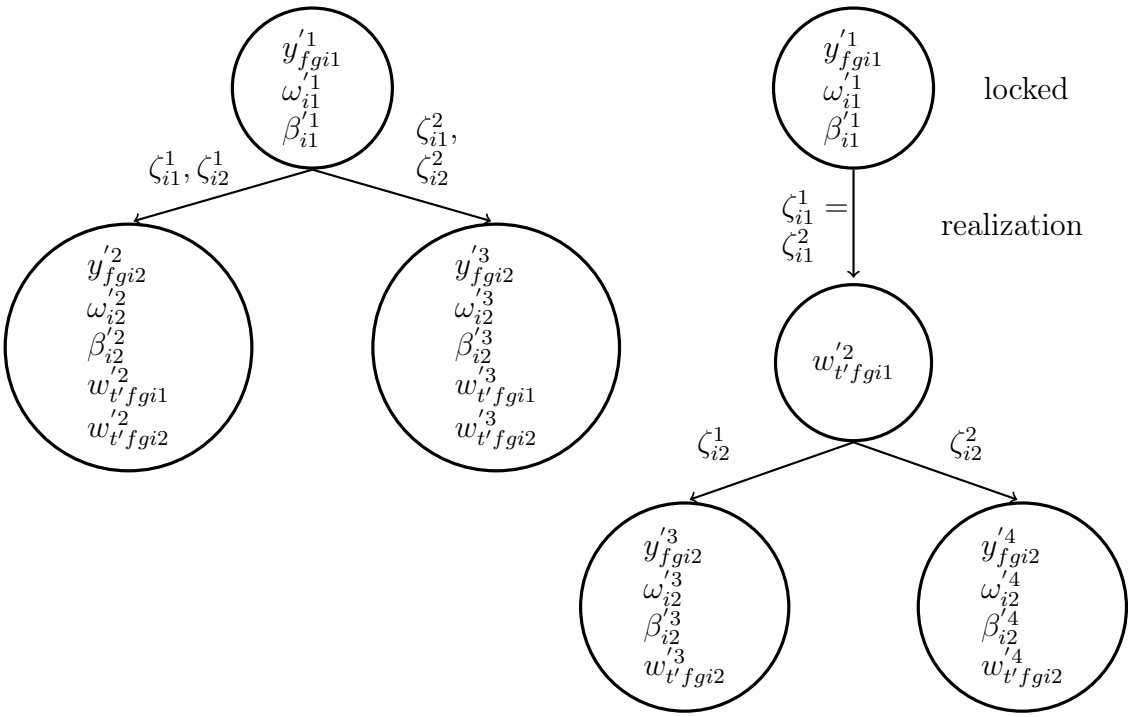

Figure 7.3: Depiction of scenario trees used to lock to the first-stage variables on the left and second-stage variables on the right.

7.2.3 Multiple Evaluations

Applying the RHH scheme once to a candidate solution gives us an evaluation of the candidate solution based on the realizations and scenario trees used. As these two factors are sampled, we would naturally like to see how our candidate solution fare under different realizations and scenario trees. Thus, we apply the RHH scheme multiple times to the same candidate solution, each application yielding an operational plan for the planning horizon together with an objective value. We let the average of the obtained objective values serve as a proxy for the expected objective value of the candidate solution, and thus its quality.

Figure 7.4 shows an application of the evaluation scheme to a candidate solution with three evaluations. The candidate solution results from a problem with a planning horizon of five years, a central period of one year, and a

Figure 7.4: Schematic overview of the evaluation of a candidate solution to the first-stage variables of the first primary period.

forecast period of four years. A central period length of one year means we use planning blocks spanning a year, and the forecast period length simply means that we look four years ahead after the end of the planning block we are currently solving. We also need to include the time periods necessary for the end of horizon constraints following the forecast period, however, we do not depict them in the figure. During the first planning block, we have graphically split the block in two to depict the difference between locking the first and second-stage variables of the planning block, while we have omitted this distinction during the later planning blocks for the sake of brevity. The graphical split of the first year does not mean that the second-stage variables follow the first-stage variables in time.

As we generate candidate solutions by using a two-stage program with a specific initial scenario tree containing a limited amount of scenarios, it might be the case that we could get a better candidate solution if we use a different initial scenario tree. We thus sample several initial scenario trees and generate a candidate solution for each tree. Each candidate solution is then run through the evaluation scheme presented above and in Figure 7.4. In the end

we have multiple candidate solutions, each with multiple evaluations.

To disentangle the impact of different stochastic realizations from the quality of the candidate solution, we use the same set of scenario trees when evaluating two different candidate solutions. The two candidate solutions are generated based on two independently sampled scenario trees, however, the scenario trees used during the evaluation of the two candidate solutions must be equal. In the example given in Figure 7.4, one scenario tree is needed to generate the candidate solution, while nine scenario trees are needed to evaluate it once, whereof five are scenario trees with a planning block realization used to lock the second-stage variables, and four are used to lock the first-stage variables. To perform three evaluations we thus need $3 \cdot 9 = 27$ scenario trees. When we evaluate multiple candidate solutions, we thus sample a total of three sets of scenario trees, each consisting nine scenario trees respecting the properties outlined in section 7.2.2 with regards to realization of the central period and conditioning on previous realizations. When evaluating two different candidate solutions, we use the first set of scenario trees during evaluation 1 of both candidate solution, and the second sampled set during evaluation 2 of both candidate solutions, and so on.

We note that the use of out of sample simulation and iterative scenario tree generation in practice renders the resulting two-stage programs infeasible. In such cases we can either return an objective value of $-\infty$ or some other appropriate value, or we can attempt to modify the program to make it feasible. Since our programs in practice always turn out to be infeasible we choose the latter option. Noting that small deviations from an operational plan with respect to timing of harvesting and the harvested cohorts' mean weight are generally acceptable, both from a business and regulatory standpoint.

7.3 Feasibility When Using Rolling Horizon Heuristics

A problem that can arise when we lock variables and generate a new scenario tree out of sample is that the new program might turn out infeasible. Locking the first-stage variables when solving for the second-stage variables of the planning block is a further complication. To illustrate this issue, assume

we lock a deployment at time t' and lock harvest of the deployment at time t. At time t the cohort has reached a mean weight of 4,001 grams, 1 gram above the weight threshold for harvesting at 4,000 grams. Moreover the total biomass of the cohort is approximately equal to the location MAB. Next, we generate a new scenario tree out of sample, and now the cohort only reaches 3,999 grams at time t, and can thus not be harvested even though it must be since we have locked in the harvest operation. Even worse, if we do not harvest, then at time $t+1$ the total biomass of the cohort may have surpassed the location MAB. Clearly our program is infeasible.

We solve the illustrated problem by moving time period t from the cohort's growth phase to its harvest phase. However, it is also possible that the cohort grows faster in one of the newly sampled scenarios, and as such must be harvested at time $t-1$ to avoid breaching the location MAB at time t. In such a case we move $t-1$ and all subsequent growth phase time periods to the cohort's harvest phase for the scenario in question. When solving the two-stage program to fix the second stage variables of the planning block, we must consider the locked variables. Denoting the locked variables with a superscript l and specifically the locked binary harvest variables as $\omega^l_{i(t-1)}$, we flip $\omega^l_{i(t-1)}$ to one to allow harvesting at time $t-1$ in the case above. To avoid forcing the company to also use its old locked harvest operations ω^l_{it}, which in itself could give rise to new feasibility problems, we only require the following initial conditions for the time periods, $\mathcal{T}^{\mathcal{C}}$, of the central period in question

$$\omega^s_{it} \leq \omega^l_{it}, \quad i \in \mathcal{I}, t \in \mathcal{T}^{\mathcal{C}}, \tag{7.1}$$

and as a consequence to allow fallowing after the last harvesting operation

$$\beta^s_{it} \leq \beta^l_{it}, \quad i \in \mathcal{I}, t \in \mathcal{T}^{\mathcal{C}}. \tag{7.2}$$

At each location we check the time periods in chronological order from $t = 1$ to $t = |\mathcal{T}^{\mathcal{C}}| + \Delta$ the latest possible time period to harvest a locked deployment. During each time period we control that the total biomass in its growth phase during $t-1$ does not exceed the location MAB at time t. When solving with locked variables we must also consider cohorts in their harvest phase, where the locked harvest variables ω^l do not allow the deployed cohorts to be harvested at time $t-1$ or earlier. If the biomass at a location respecting the

aforementioned characteristics does exceed the location MAB at time t, we move $t-1$ and the cohort's subsequent growth time periods to the cohort's harvest phase. When solving the two-stage program to fix the second stage variables we set $\omega^l_{i(t-1)} = 1$ to allow harvesting. When $t-1 > |\mathcal{T}^{\mathrm{C}}|$ we do not consider $\omega^l_{i(t-1)}$ as it does not exist, and when we solve the two-stage program to lock in the first-stage variables ω^l does not exist at all. In the case were a deployment consists of several cohorts we perform the above procedure on the cohort with the highest mean weight. Next, we check whether the time period has become feasible, if not we move onto the second heaviest cohort in the deployment, and so on.

After ensuring feasibility at a location level, we perform the same procedure at a company level with respect to the company wide MAB. We perform the location level and company wide level feasibility procedure described above for each scenario. As a final step we ensure that if a cohort's harvest phase starts during a time period t' which is contained in the central period in one of the scenarios, that the cohort's harvest phase starts no later than during time period t' in all of the scenarios. We achieve this by simply moving time periods from the growth phase to the harvest phase in the scenarios where it is necessary. The reasoning is that if we harvest in one of the scenarios, we must harvest in all due to the non-anticipativity constraints.

As the astute reader might have noted, constraints limiting harvest volumes are likely to cause infeasibility when solving with locked variables or initial conditions. We therefore relax the harvest volumes constraints, and penalize breaching them trough adding penalty terms to the objective function. In our case only the company-wide upper harvest constraint (6.8) and the lower bound location level harvest constraint (6.7) need to be relaxed. We set the upper bound of the location harvest constraint equal to the location MAB in our case studies, making it non-restrictive. We thus omit relaxing it here for the sake of brevity, however, for different upper bounds it might be necessary to relax it. For simplicity we relax the constraints for all time periods, although it suffices to relax the constraints for time periods $1, 2, \ldots, |\mathcal{T}^{\mathrm{C}}| + \Delta$.

We replace the location lower harvest volume constraints by

$$\sum_{f \in \mathcal{F}} \sum_{g \in \mathcal{G}} \sum_{t' \in \mathcal{T}^{\mathrm{D}}_{fgit}} w^s_{t'fgit} \geq \mathrm{L}^{\mathrm{H}}_i \omega^s_{it} - \tilde{\phi}^s_{it} \quad s \in \mathcal{S}, i \in \mathcal{I}, t \in \mathcal{T}, \tag{7.3}$$

where $\tilde{\phi}^s_{it}$ is a continious decision variable with its domain restricted to lie between zero and the location's lower harvest volume limit $\mathrm{L}^{\mathrm{H}}_i$ for tightness purposes

$$0 \leq \tilde{\phi}^s_{it} \leq \mathrm{L}^{\mathrm{H}}_i \omega^s_{it} \quad s \in \mathcal{S}, i \in \mathcal{I}, t \in \mathcal{T}. \tag{7.4}$$

Similarly, we replace the company wide upper harvest constraint by

$$\sum_{f \in \mathcal{F}} \sum_{g \in \mathcal{G}} \sum_{i \in \mathcal{I}} \sum_{t' \in \mathcal{T}^{\mathrm{D}}_{fgit}} w^s_{t'fgit} \leq \mathrm{U}^{\mathrm{COMP}}_t + \hat{\phi}^s_t \quad s \in \mathcal{S}, t \in \mathcal{T}, \tag{7.5}$$

where we restrict the decision variable $\hat{\phi}^s_t$ to be non-negative and simply yield an upper harvest constraint that potentially can be as relaxed as the company wide MAB

$$0 \leq \hat{\phi}^s_t \leq \mathrm{MAB}^{\mathrm{COMP}} - \mathrm{U}^{\mathrm{COMP}}_t \quad s \in \mathcal{S}, t \in \mathcal{T}. \tag{7.6}$$

We then simply subtract the following penalty term from our objective, to discourage the model from utilizing these relaxations unless it is necessary

$$\sum_{s \in \mathcal{S}} \pi^s \left[\sum_{t \in \mathcal{T}} \left[\hat{\phi}^s_t + \sum_{i \in \mathcal{I}} \tilde{\phi}^s_{it} \right] \right]. \tag{7.7}$$

We note that the procedures outlined above do far from guarantee feasibility. Especially the activity constraints that require the company to deploy biomass in combination with the lower bound constraints on deployment volumes may still cause infeasibility. This occurs when the company is forced to deploy fish, but are unable to do so without breaching the company wide MAB due to the lower bound on deployment volumes. We choose to not relax the deployment lower bounds in the same manner we did with harvest volumes, as we find that the model could be forced to deploy very small deployments consisting of less than a net-pen worth of fish, which is not feasible from a practical standpoint. As such, we decide to leave the RHH framework infeasible if this happens.

Chapter 8

Case Study

In this chapter we introduce the case study which forms the basis for the computational study we present in chapter 9. We start of by giving a brief introduction to Eidsfjord Sjøfarm, which is an Atlantic salmon farming company situated in Northern Norway, and serves as the inspiration for our case study. Thereafter we introduce the characteristics of the relevant smolt types. Next, we specify the parameters needed to model the presence of salmon lice treatments and present a complete scenario generation algorithm. Lastly we finish of by presenting the remaining constraint parameters and the reasoning behind their chosen values.

8.1 Company

Eidsfjord Sjøfarm is a mid-size Norwegian Atlantic salmon farming company situated in Northern Norway, which historically has produced approximately 15,000 head-on-gutted tonnes of Atlantic salmon per year. For the sake of our case study we base our production system on Eidsfjord Sjøfarm's historic production system, which consisted of 16 locations scattered across Northern Norway as presented in Figure 8.1. Table 8.1 gives the location MAB for each of the 16 locations, together with the initial biomass present at the start of our case study, which is November 2020. Ten of the locations are located in Vesterålen, three are situated around Senja, and the final three locations are located in Nord-Troms. The regions lie within the regulatory defined production zones 9, 10, and 11 respectively. However, Eidsfjord Sjøfarm has been granted an inter-regional MAB between these three zones, yielding a

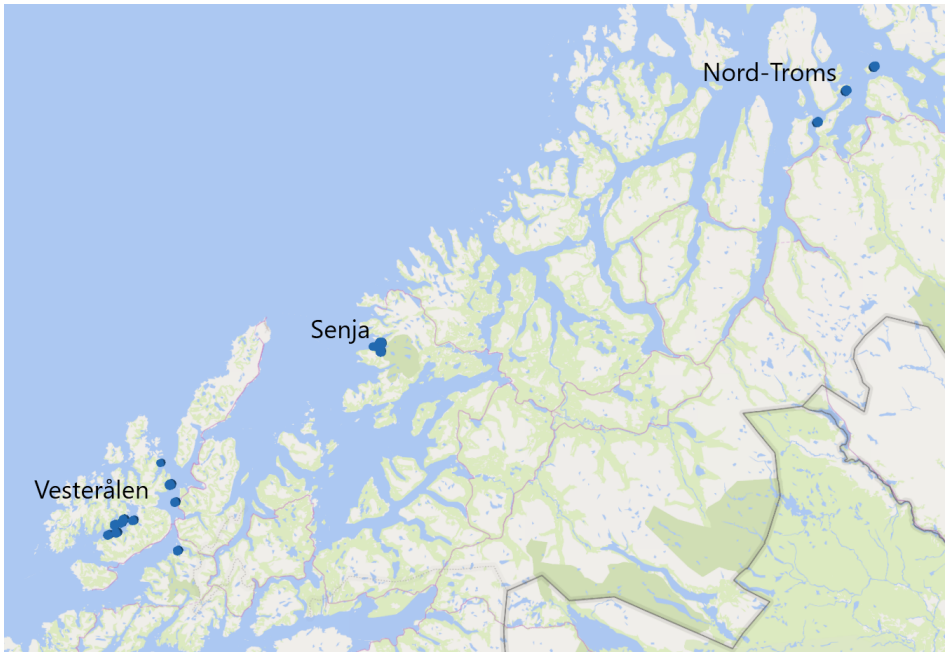

Figure 8.1: Geographical location of Eidsfjord Sjøfarm's farms.

company wide MAB of 10,902 tonnes. We include the fact that Eidsfjord Sjøfarm's current production system consists of 17 locations and a company wide MAB of 14,528 tonnes, however, we choose to use the historic production system as historical benchmark data are available for this system, providing us the possibility to sanity check our results.

8.2 Smolt and Salmon Specifications

We now define the parameters related to smolt and salmon characteristics, which includes release weights, TGC and mortality rates of the different smolt types, harvest weight requirements, and specification of cohorts' weight distributions.

8.2.1 Release Weights

The company has the opportunity to release smolt during eight months of the year, which are split into three main release windows. The initial weight

Table 8.1: Overview of Eidsfjord Sjøfarm's production system with initial conditions.

ID	Name	Region	MAB	Initial Biomass (tonnes)	Initial Mean Weight (grams)	Months at Sea
1	Bremnesøya	Vesterålen	3900	1864	2207	10
2	Daljorda	Vesterålen	3120	2316	3273	13
3	Innerbrokløysa	Vesterålen	3120	0	0	NA
4	Kuneset	Vesterålen	3120	0	0	NA
5	Langholmen N	Vesterålen	3120	0	0	NA
6	Pollneset	Vesterålen	3120	0	0	NA
7	Reinsnesøya	Vesterålen	3120	1007	1464	5
8	Sandan Sø	Vesterålen	2340	0	0	NA
9	Stretarneset	Vesterålen	3120	1033	5261	17
10	Trolløya Sv	Vesterålen	3120	999	2641	5
11	Flesen	Senja	2700	0	0	NA
12	Kvenbukta V	Senja	2700	225	317	2
13	Lavika	Senja	2700	571	689	2
14	Hagebergan	Nord-Troms	3600	1497	6000	17
15	Haukøya Ø	Nord-Troms	3600	0	0	NA
16	Russelva	Nord-Troms	3500	1298	1591	5

of deployed smolt depends on which release window they are deployed in, as presented in Table 8.2. During the release window containing December and January the company must use so called winter smolt with an initial weight of 250 grams. A further constraint during this release window is that the company can only release 550,000 smolt due to a supply limitation on winter smolt. We formally have $\mathrm{S}_p = 550$ during each period $p \in \mathrm{P} = \{\{\text{December Year } y, \text{January Year } y\} : y \in \mathcal{Y}\}$, where $\mathcal{Y}$ is the set of years in the planning horizon. During the release window of April, May, and June the initial release weight is constrained to 150 grams, while the release window containing the months of July, August, and September utilizes an initial smolt weight of 100 grams. In general the initial smolt weight is higher when sea temperatures are low, as more robust smolt are required.

We add that in the industry the weight used during the summer half re-

Table 8.2: The possible release months for each initial smolt weight.

Weight	Months
100 grams	July, August, September
150 grams	April, May, June
250 grams	December, January

lease windows are no longer as strictly defined as it traditionally has been. Recent trends show a tendency towards deploying smolt in the weight range 100 – 200 grams during both of the summer half release windows, and the trend is more and more towards the higher end of this interval. However, we have kept a more traditional deployment weight scheme for the sake of simplicity.

8.2.2 Smolt Types

We study five smolt types, as introduced in section 2.4, namely, regular smolt, mono-female smolt, mono-male smolt, smolt originating from cooled eggs, and early gender maturation smolt. The choice of these five smolt types stems from the fact that Aquagen is currently running studies on them.

All smolt types except the early maturation type can be released during any of the release months presented in Table 8.2. However, the early maturation smolt can only be released during June and July, since its growth boost can only occur during July, August, and September, and we require that both its 13th and 14th month at sea occur during the three aforementioned months.

The five smolt types differ with respect to TGC and mortality rates. When estimating these parameters we only have reliable sea trial data available for the regular, mono-female and mono-male smolt types. Figure 8.2 presents the TGC with respect to time spent at sea for the five smolt types. As apparent, the cooled eggs and mono-male smolt types exhibit the highest growth rates, save for the period when early maturation experiences its growth boost, while mono-female is the slowest growing fish. The TGC of the regular, mono-female, and mono-male smolt types are based on data from sea trials that lasted 14 months. To fill in the blanks during month 15 and later, we simply extrapolate by using the average of months 12 – 14. For the early maturation smolt type we assume that it has to be harvested during it 13th or 14th month

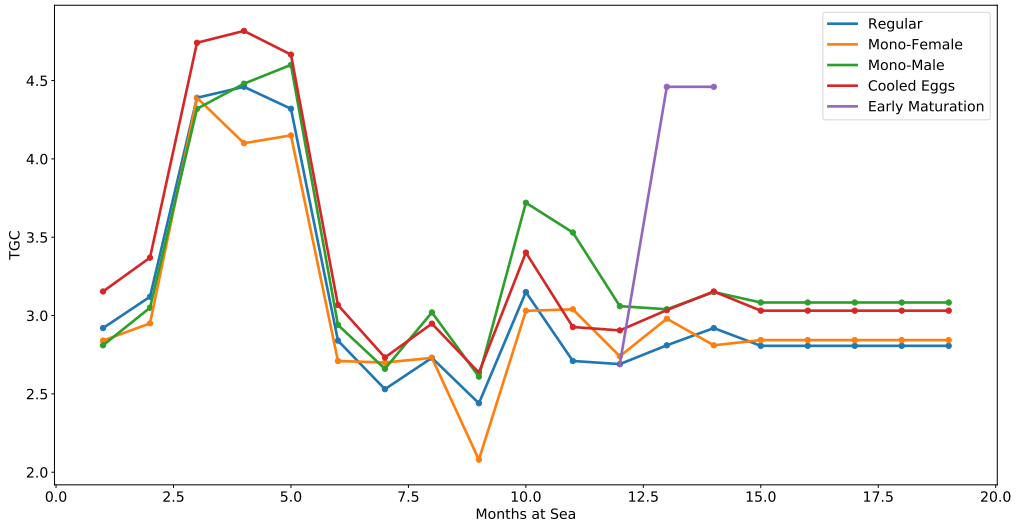

Figure 8.2: TGC of the different smolt types as a function of time spent at sea.

at sea, and that it gets a growth boost during these months. We assume that the early maturation smolt type's TGC is equal to the regular smolt type until month 13 and 14 where it becomes equal to the highest TGC value the regular smolt type exhibited, i.e. the regular smolt type's TGC value during its 4th month at sea. For the cooled eggs we assume that its TGC is 8 % higher than regular smolt type's TGC, which is based on fresh water trials, where the cooled eggs exhibited an 8 % higher TGC than regular smolt. We do note, however, that TGC is not really applicable during the fresh water stage, but for the lack of any better data we choose to continue with this assumption.

As for mortality, we model the mortality rates as a function of months enjoyed at sea. We base the mortality rates for the regular, mono-female, and mono-male smolt types on empirical values obtained from sea trials. We assume that cooled eggs and early maturation smolt experience the same mortality rates as the regular smolt type. We do however, scale the mortality rates such that the regular smolt type match the national average and the ratios between the different smolt types remain unchanged. We do this through replacing the regular smolt types mortality rate with the national average

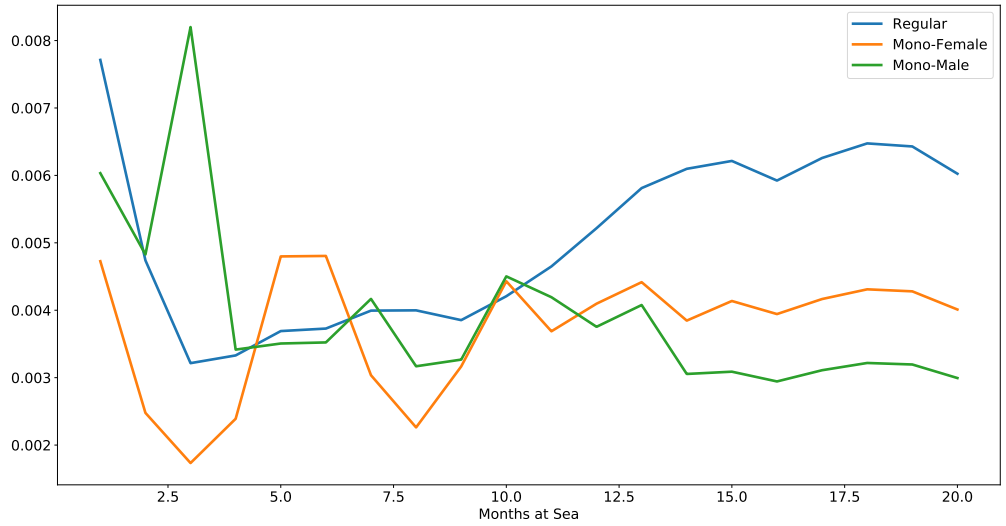

Figure 8.3: Scaled mortality rates for the regular, mono-female, and mono-male smolt type. We assume that the cooled eggs and early maturation smolt types exhibit the same mortality rates as the regular smolt type.

obtained from Bang Jensen et al. (2020), and then scaling the other smolt types mortality rates such that the ratios from the sea-trails are preserved. As with TGC the ratios for months 1 through 14 are based on the sea trial data, while the months $15-19$ are based on the average during the months 12 through 14. The early maturation smolt type's mortality rate is naturally cut short after 14 months, its latest possible harvest time. Figure 8.3 gives a graphical depiction of the monthly mortality rates for the regular, mono-female, and mono-male smolt type as a function of months enjoyed at sea. An interesting point to examine is the impact of the mono-female smolt type's low mortality rate during its first months at sea, and whether this compensates for its comparatively low growth rates.

8.2.3 Other Fish Paramters

We now provide the missing parameters with regards to fish modeling. We assume that the HOG yield of a harvested salmon is $\kappa = 0.84$ based on Mowi ASA (2020). We define the harvest threshold as 3,500 grams HOG, giving us a harvest threshold in terms of sea weight of $\frac{3500}{0.84} \approx 4,166.67$

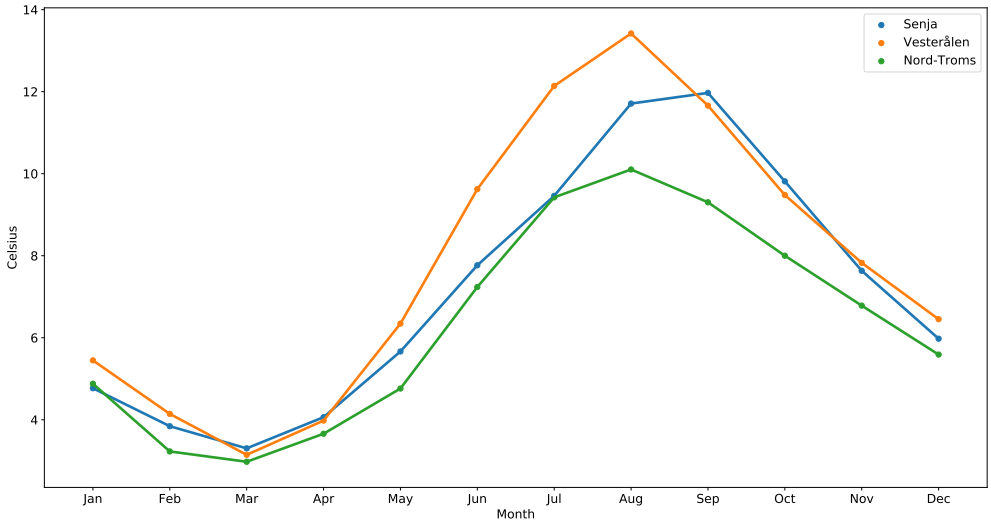

Figure 8.4: Temperature curves in Celsius for the three regions as a function of calendar month.

grams. Hence in terms of our modeling language, a cohort with a mean weight below this threshold is in its growths phase, while a cohort with a mean weight at or above this threshold is in its harvest phase. Similarly, we define an upper harvest weight threshold of $6,500$ HOG grams, i.e. cohorts must be harvested before their mean sea weight reaches $\frac{6,500}{0.84} \approx 7,738.10$ grams, meaning a cohort's harvest time periods set does not contain time periods where the cohort's mean weight would be above this threshold. These thresholds are based on the fact that market demand is considerably lower for weights outside of the range defined by the lower and upper thresholds. Moreover, we assume that a cohort's weight distribution has a coefficient of variation $\text{CV} = 0.225$ based on Forsberg (1999) and Hæreid (2011).

8.3 Temperatures

To predict weight growth using TGC, we need to specify the temperature at each location for all time periods. We assume that all locations situated within the same region have the same temperature. As previously mentioned, Eidsfjord Sjøfarm's locations are scattered across three regions, which from south to north are denoted Vesterålen, Senja, and Nord-Troms. Figure 8.4

gives the average monthly seawater temperatures for these regions based on historic data obtained from Eidsfjord Sjøfarm. Senja and Vesterålen's curves are based on historic data from 2019. Nord-Troms lacked temperature data during the months January to April in 2019, we therefore impute these values. We estimate the temperature for these months assuming the temperature ratio between the Nord-Troms and Senja was the same in 2019 and 2018. We then compute this ratio for each of these months in 2018 based on available data, and use these ratios to impute the missing values for Nord-Troms in 2019. We assume that the temperature repeats yearly. Thus, we obtain the temperature curves for a multi-year planning horizon by duplicating Figure 8.4 for each year of the planning horizon and the time periods contained in the end of horizon conditions.

8.4 Salmon Lice Parameters and Scenario Generation

We now provide a complete scenario generation algorithm for salmon lice treatments based on the approach presented in section 5.2.

8.4.1 Salmon Lice Treatment Data

This subsection defines the parameters that are needed before we can perform scenario generation. The production locations introduced in section 8.1 are split between three production zones as defined by Norwegian authorities, namely, production zones 9, 10, and 11. As such we define the lice treatment probability parameters $\hat{\rho}_{it} = \mathrm{P}(\zeta_{it} = 1|\zeta_{i(t-1)} = 0)$ and $\check{\rho} = \mathrm{P}(\zeta_{it} = 1|\zeta_{i(t-1)} = 1)$ separately for these three zones, since lice data are available at a production zone level. Figure 8.5 shows the values of $\hat{\rho}$ and $\check{\rho}$ as a function of calendar month for production zones 9, 10, and 11 respectively. The general trend is that production zone 11 has considerably lower probability of lice treatments than zones 9 and 10, especially during summer.

When a treatment is performed we assume that five feeding days are lost at the location, i.e. $\mathrm{D} = 5$. Lastly, based on Overton et al. (2019) we assume mortality rates increase by one percentage point during the time period when treatment occurs at a location, yielding $\hat{\mu}_{it} = 0.01$.

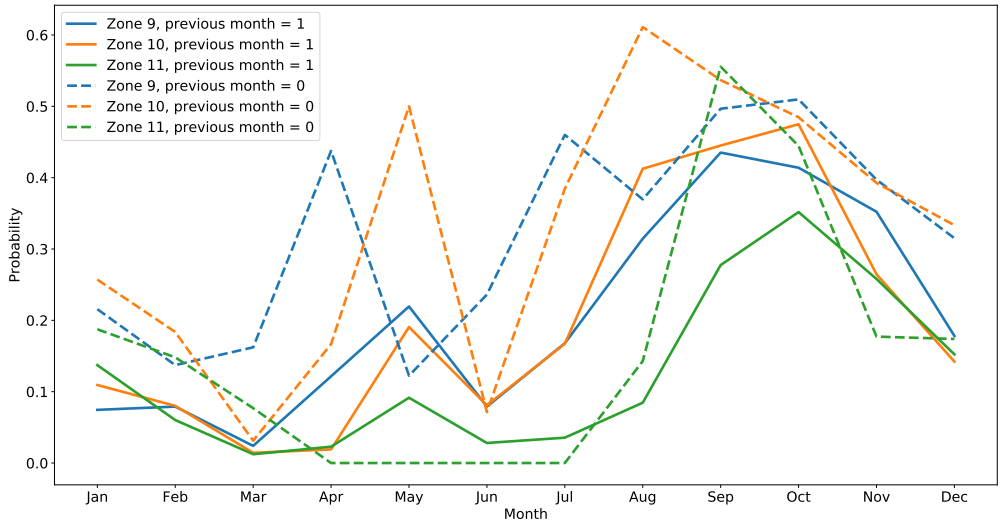

Figure 8.5: The parameters $\hat{\rho}$ and $\check{\rho}$ for production zones 9, 10, and 11 as a function of calendar month. Previous month = 0 refers to $\hat{\rho}$, while previous month = 1 refers to $\check{\rho}$.

8.4.2 Scenario Generation

We start off with deciding the parameters used to model the relative risk which depends on the seaway distance d_{ij} in kilometers between location i and j, and is given as

$$\mathrm{RR}_{ij} = \max\{\mathrm{N} \cdot e^{-\frac{\phi_1}{\phi_2} d_{ij}^{\phi_2}}, 1\}. \tag{8.1}$$

Aldrin et al. (2017) estimate that $\phi_1 \approx 0.351$ and $\phi_2 \approx 0.568$. We decide that two locations are pairwise independent of one another if they have a seaway distance of at least 50 km between them, hence we define $\mathrm{N} = e^{\frac{0.351}{0.568} \cdot 50^{0.568}}$. We measure the seaway distance manually by using the distance measuring tool in Google Maps, the results of these measurements are given in Appendix A. As a consequence of our chosen value of N we can generate scenarios for each of the three production zones separately as they are uncorrelated. This reduces the computational burden significantly, as the largest sample space we have to estimate probabilities for is of size $2^{10} = 1024$ instead of $2^{16} = 65,536$ for each time step. We still have to consider two extra sample spaces, each of size $2^3 = 8$ separately, but the computational overhead of this is

minuscule. The scenario generation algorithm based on the presentation in section 5.2 is given in algorithm 1. We assume that $\zeta^s_{i0} = 0$ at the start of the planning horizon. Lastly, we assign all sampled scenarios the same probability, i.e. $\pi^s = \frac{1}{|\mathcal{S}|}$.

Algorithm 1 Scenario generation for a production region $\mathcal{I}' \subseteq \mathcal{I}$ given initial conditions ζ_0, relative risks RR_{ij}, $\hat{p}_t$, $\check{p}_t$, and the sample space $\Omega_{\mathcal{I}'} = \{0,1\}^{|\mathcal{I}'|}$.

for all $t \leftarrow 1, 2, \ldots, |\mathcal{T}|$ **do**
  $\rho_t \leftarrow \hat{\rho}_t \odot (1 - \zeta_{(t-1)}) + \check{\rho}_t \odot \zeta_{(t-1)}$

  // Find the corrected joint distribution between pairwise locations.
  $\mathrm{P}'^t_{ij} \leftarrow 0$
  for all $(i,j) \in \mathcal{I}' \times \mathcal{I}'$ **do**
    $\mathrm{P}^t_{ij} \leftarrow \frac{\rho_{it}\rho_{jt}}{1+\rho_{jt}(\mathrm{RR}_{ij}-1)}\mathrm{RR}_{ij}$
    $\mathrm{P}'^t_{ij} \leftarrow \min\{\mathrm{P}^t_{ij}, \mathrm{P}^t_{ji}\}$
  end for
  Use iterative proportional fitting to find $\mathrm{P}(\zeta_t = c)$, respecting
  $\mathrm{P}(\zeta_t = \mathbf{1}_n) = \rho_t$ and $\mathrm{P}(\zeta_{it} = 1 \wedge \zeta_{jt} = 1) = \mathrm{P}'^t_{ij}$

  // Sampling ζ_t from a discrete distribution
  Sample $z \in \mathcal{U}(0,1)$
  $r \leftarrow 0$
  for all $c \in \Omega_{\mathcal{I}'}$ **do**
    $r \leftarrow r + \mathrm{P}(\zeta_t = c)$
    if $r \geq z$ **then**
      $\zeta_t \leftarrow c$
      break
    end if
  end for
end for
return ζ

8.5 Other Constraint Parameters

This section summarizes the last parameters we need to specify, an overview of these parameters is provided in Table 8.3.

We set the maximum rearing cycle length Δ to 19, which is done through simply adjusting the cohorts' harvest time periods sets to ensure that they are harvested before Δ time periods have passed. The Norwegian regulations stipulate that a site must be fallowed for at least two months between rearing cycles, hence we have $\Lambda = 2$. In a similar fashion, the regulations require the company to have some biomass present at each location at least once during a 24 month period yielding $\Gamma = 24$.

We let $\mathrm{Q}_i = 4$, giving the lower bound on the amount of biomass that must be employed at a location if biomass is employed. This value is somewhat arbitrary, and is set to a low non-zero value to simply enforce the relationship that $\beta_{it} = 1$ if and only if biomass is employed at location i at time t, without being restrictive in practice. We want to avoid a tight lower bound, as it causes feasibility issues when applying our RHH framework. The value of 4 is chosen to be absolutely certain the solver does not encounter any numerical issues, which could occur if we choose an overly low value.

We base our deployment restrictions on the number of fish that fit in a net pen, which in practice is around 120 thousand fish. We restrict the lower deployment volumes of a cohort to two net pens, i.e. $\tilde{\mathrm{L}}_i^{\mathrm{DEP}} = 240$ thousand fish. At a location level we demand the deployment volumes to lie in the range between $\hat{\mathrm{L}}_i^{\mathrm{DEP}} = 480$ and $\hat{\mathrm{U}}_i^{\mathrm{DEP}} = 1,440$ thousand fish, which is between 4 and 12 net pens. As we allow a deployment to consist of only one cohort we let $\tilde{\mathrm{U}}_i^{\mathrm{DEP}} = \hat{\mathrm{U}}_i^{\mathrm{DEP}}$.

Our harvest volume limits are motivated by the fact that Eidsfjord tries to harvest at least $1/2$ of a net pen per week, which is equivalent to approximately 475 tonnes. Note that we measure harvest volumes in tonnes and not in the number of fish. We set the lower harvest limit at a location level per time period to $\mathrm{L}_i^{\mathrm{HARV}} = 425$, which is equivalent to 50 tonnes less than $1/2$ of a net pen at full capacity. We limit the harvest volumes at a company level to be less than $\mathrm{U}_t^{\mathrm{COMP}} = 3,800$, which is approximately twice the monthly harvest goal. We do not restrict the harvest volumes at a location level any further, as such we set $\mathrm{U}_i^{\mathrm{HARV}} = \min\{\mathrm{MAB}_i, 3800\}$

Table 8.3: Summary of parameters regarding production cycles length, fallowing length, activity constraints, deployment volume constraints, and harvest volume constraints.

Parameter	Value
Δ	19
Λ	2
Γ	24
Q_i	4
$\tilde{\mathrm{L}}_i^{\mathrm{DEP}}$	240
$\tilde{\mathrm{U}}_i^{\mathrm{DEP}}$	1,440
$\hat{\mathrm{L}}_i^{\mathrm{DEP}}$	480
$\hat{\mathrm{U}}_i^{\mathrm{DEP}}$	1,440
$\mathrm{L}_i^{\mathrm{HARV}}$	425
$\mathrm{U}_t^{\mathrm{COMP}}$	3,800
$\mathrm{U}_i^{\mathrm{HARV}}$	$\min\{\mathrm{MAB}_i, 3800\}$

Chapter 9

Computational Study

In this chapter we perform a computational study based on the case study introduced in chapter 8. All experiments are performed on a linux computational cluster with a Lenovo NextScale nx360 M5 server with a 2x 2.3GHz Intel E5-2670v3 – 12 core CPU, 64GB RAM, and a 120GB SATA SSD. The software used is Python 3.7.4, R 3.6.2, and Xpress 8.10.

9.1 General Configurations

We now provide the general setup for the computational instances we study in this chapter. The instances are based on chapter 8, however, we now provide the details for the specific instances we study, together with details on the technical configurations we use for the MILP-solver and our rolling horizon heuristic framework. All instances we study in this chapter have a planning horizon of five years and each time period represents one calendar month.

9.1.1 Smolt Combinations

We study three different smolt combinations that the farming company has available. The company has the opportunity to use all smolt types she has available, but she does not need to use all of them. The reason for studying different combinations is to measure the potential value of introducing different novel smolt types, e.g. how much of an increase in objective value can the company achieve by introducing mono-gender smolt compared to only

using the regular smolt type.

The first combination we study is the regular smolt type in isolation, the second is regular smolt, mono-female, and mono-male smolt, and the last combination includes all smolt types, i.e. cooled eggs and early maturation smolt in addition to the ones included in the second combination. The second combination is included since it is based on more reliable data than the third combination.

9.1.2 Rolling Horizon Heuristic

The rolling horizon heuristic (RHH) framework we introduced in chapter 7 is applied with a central period of one year and a forecast period of four years. Meaning, we use planning blocks spanning a year, and for each subproblem we look four years ahead after the end of the planning block we are solving. We thus consider a time span of the same length as the five year planning horizon for each subproblem. However, the forecast period extends beyond the planning horizon during later iterations, e.g. during the fifth main iteration we consider the years five through nine. Although the central and forecast period defining the two-stage stochastic program contain 60 time periods, we ensure that each scenario is generated to cover 80 time periods, to cater to the end of horizon time periods following the forecast period. We also relax integrality requirements during the last two years of the forecast period. Each time we solve the two-stage stochastic program we let the program run until an optimality gap of 3 % is achieved or three hours have lapsed. If no feasible solution has been found within three hours, the program runs until a feasible solution has been found.

9.1.3 Evaluation Measurements

The two most obvious and important ways to compare different candidate solutions are the average objective value and average harvest volumes, where the average is computed over all evaluations of the candidate solution. However, in order to assess the industrial and technical implications of our results, it is necessary to introduce some additional measurements. We first introduce as set of performance measurements that are commonly used in the industry. These serve two purposes, one, they allow us to sanity check our results against industrial practices, second, they allow us to examine what

drives the change in objective value when we examine different smolt combinations. Lastly, we introduce a sensitivity measure to assess the difference between the best and worst outcome of a candidate solution.

Industrial Key Performance Indicators

Firstly, we compute the average HOG harvest weight as

$$\kappa \cdot \frac{\sum w_{t'fgit} \cdot \mathrm{W}_{t'fgit}}{\sum w_{t'fgit}}, \tag{9.1}$$

where $\mathrm{W}_{t'fgit}$ is the mean weight of the cohort corresponding to the harvest volume variable $w_{t'fgit}$, we omit the summation indices for the sake of brevity. Equation (9.1) yields the average harvest weight per harvested tonne of biomass, and not the average weight per harvested fish. This is a simplification we utilize, since we do not keep track of the amount of fish in a cohort during the rearing cycle. Ceteris paribus a higher average HOG harvest weight is advantageous as it increases production volumes.

We define the short rearing cycle length (SRCL) as the average number of time periods from deployment to the first harvesting operation of the deployment takes place. Similarly, we define the long rearing cycle length (LRCL) as the average number of time periods from deployment to the last harvesting operation performed on said deployment. Note that we measure based on deployments and not cohorts, e.g. if two cohorts are part of a deployment the SRCL only considers the cohort that is first subject to harvesting. In general it is preferable to achieve a given average HOG harvest weight with short rearing cycles, as this would allow the company to increase their number of rearing cycles reaching the same harvest weight during a planning horizon. When calculating the average cycle lengths we only consider deployments performed during the planning horizon, i.e. we omit the biomass present through initial conditions at the start of the planning horizon. When calculating the long rearing cycle length we only include deployments that are completely harvested at the end of the planning horizon.

Sensitivity Measurements

To measure the difference between the best and worst outcome of a candidate solution, we define the following gap measurement

$$\text{GAP}(\text{OBJ}) = \frac{\max_i \text{OBJ}_i - \min_i \text{OBJ}_i}{\min_i \text{OBJ}_i}, \tag{9.2}$$

where OBJ is a vector where each element i correspond to a measurement value of evaluation i of the candidate solution in question. Despite its similarities, this gap measurement should not be confused with an optimality gap. We introduce it, since we cannot provide an optimality gap for an entire evaluation of a candidate solution when we apply an RHH. Our gap can also be applied to other values than just the objective value, such as total harvest volumes or solve time. In general a low gap indicates that a candidate solution is stable with regards to different realizations of the future, while a high gap tells us that the candidate solutions performance heavily depends on the realization of the future.

9.2 Size of the Two-Stage Program

In this subsection we give a brief overview of the computational complexities we are faced with when studying instances varying in the number of smolt types and scenarios. Table 9.1 gives the size of the two-stage stochastic program in columns, rows, and non-zero elements. The numbers are provided after the solver has performed pre-solve, which reduces the original size of the program. The size reduction resulting from pre-solve varies somewhat depending on the precise scenarios/program, but the numbers presented in Table 9.1 are representative for the program size of other scenarios as well. The five scenario programs used in Table 9.1 share the same scenarios across the different number of smolt types, while the one scenario programs simply have $\zeta_{it}^s = 0$ for all i and t.

Naturally as we increase the number of smolt types and scenarios, the number of rows, columns, and non-zero elements increases. However, this scales sub-linearly in the case of increasing smolt types. Going from one smolt type to five in the one scenario case, we increase the number of rows and columns by approximately 3.1x and 2.7x respectively, and by 3.2x and 2.7x in the case with five scenarios. While increasing the number of scenarios increases the

Table 9.1: Size of the two-stage stochastic program in rows, columns, and non-zero elements after pre-solve for different number of available smolt types and scenarios.

Smolt Types	Scenarios	Rows	Columns	Non-Zero Elements
1	1	14,760	4,933	57,077
1	5	68,829	21,025	262,160
3	1	35,895	10,547	134,061
3	5	170,469	44,649	620,862
5	1	45,268	13,349	171,638
5	5	218,066	56,399	801,477

number of rows by 4.7x and the number of columns by 4.2x on average. The number of non-zero elements increase on average by 4.6x when going from one to five scenarios, and 3.0x when increasing the number of available smolt types from one to five. These numbers show that increasing the number of scenarios is more costly in terms of program size than increasing the number of smolt types. However, it should be noted that solution-time does not scale with the numbers presented in this paragraph. Indeed, the results we present later show that going from one to five scenarios on average increases solve-time by 25.8x. Moreover, going from one smolt type to five smolt types increases solve-time by 9.0x on average, while going from one to three smolt types increases it by 3.74x on average. This further underlines the fact that increasing scenarios is more costly in terms of solution time than increasing the number of smolt types.

9.3 Maximizing Harvest Volumes

In this section we study instances where we maximize harvest volumes. This is done through setting the relative price $P_t^\alpha = 1$ for all weight classes α and all time periods t. We start off by doing a deterministic study without lice to establish a benchmark and to evaluate the usage of the RHH. Next, we perform a stochastic study with lice treatments, and thereafter we evaluate the value of using stochastic planning.

Table 9.2: Results of the deterministic no-lice study.

Smolt Types	Solution Method	Objective Value	Optimality Gap	Harvest Volumes	Avg. HOG Harv. Weight	Solution Time [s]
Reg.	No Heuristic	88,969.65	0.02 %	88,969.65	4,680.03	10,800
Reg.	RHH Relax 0y	88,092.70		88,092.70	3,975.69	263
Reg.	RHH Relax 2y	88,822.93		88,829.77	4,061.93	228
Reg. & Mono	No Heuristic	96,783.58	0.17 %	96,783.58	4,783.89	10,800
Reg. & Mono	RHH Relax 0y	96,661.10		96,661.10	3,998.85	1,119
Reg. & Mono	RHH Relax 2y	95,872.13		95,945.17	4,085.41	679
All	No Heuristic	97,135.91	0.63 %	97,135.91	4,593.56	10,800
All	RHH Relax 0y	97,458.93		97,458.93	3,853.47	2,113
All	RHH Relax 2y	96,774.69		96,774.69	3,883.68	1,559

9.3.1 No-Lice Study

To assess the effect of salmon lice we study a deterministic instance without lice to establish a no-lice benchmark. The deterministic instance is defined by one scenario with all $\zeta_{it}^1 = 0$, i.e. no lice treatments are necessary. We also use the opportunity to measure the effect of applying the RHH in a non-stochastic setting. We compare the result of the RHH to the result of solving the same deterministic instance without applying any heuristic, which is the same as solving the program given in chapter 6 with one scenario. When we do not apply RHH, we simply let the solver run for three hours, or until a solution is found if no solution is found within three hours. Lastly, we examine the effect of relaxing the integrality requirements during the later time periods of the forecasting period during iterations of the RHH.

The main results of the deterministic study are presented in Table 9.2, where relax 2y refers to relaxing the integrality requirements during the last two years of the forecast period in the RHH. We keep our presentation of the results at a high level, as the main purpose of the deterministic no-lice study is to measure the effect of different versions of the RHH without the presence of stochastic elements, and to provide a benchmark of production levels in a no-lice setting.

To our surprise, the RHH does not impact the objective value or harvest volumes to any significant degree. For the all smolt types combination, the objective value and harvest volumes are actually higher when the RHH Relax 0y is applied than when no heuristic is applied. The major difference we see, is that the average harvest weight measured in grams is considerably lower

when the RHH is applied, however, as discussed this does not have any major impact on the objective value or total harvest volumes.

Additionally, relaxing integrality during the final two years of the forecasting period yields a modest reduction in objective value and harvest volumes. Indeed, when we have only the regular smolt type available we even see a slight increase in both values. We conclude, that we can safely utilize an RHH where we relax the integrality requirements during the last two years of the planning horizon, without reducing the quality of our solutions to any significant degree, while saving compute time. Another interesting aspect is that the harvest volumes are equal to the objective value in almost all cases, meaning that the relaxations of harvest volume constraints as explained in subsection 7.3 are rarely exploited.

9.3.2 Stochastic Study

We now perform a stochastic study utilizing stochastic programming to perform production planning. The stochastic instances have five scenarios per two-stage program, where the scenarios are sampled according to Algorithm 1 given in section 8.4. We evaluate three candidate solutions as explained in section 7.2 for each smolt combination, each candidate solution is generated by using an individually sampled scenario tree. Each candidate solution is then evaluated 15 times, as such we sample 15 sets of nine scenario trees to use during the evaluations. This ensures that the different candidate solutions are evaluated using the same sets of scenario trees, e.g. the third evaluation of the first candidate solution uses the same scenario trees as the third evaluation of the second candidate solution, and so on as we explained in the last paragraph of section 7.2. Hence, the evaluations measure the differences contained in the candidate solutions in isolation.

The average number of lice treatments per location per year for the 15 different realizations of the planning horizons is 2.41, while the minimum and maximum are 1.98 and 2.80 respectively. Figure 9.1 visualizes the realization of the stochastic lice treatment variables ζ_{it}^s per location and time period for the first realization we sampled, green indicates no lice treatment being necessary and red indicates that a lice treatment is necessary. Note the strong correlation between locations in the same production zone, i.e. locations $1-10$, $11-13$, and $14-16$.

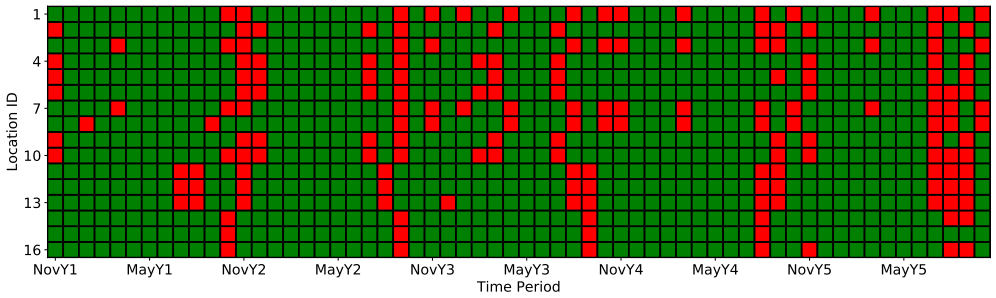

Figure 9.1: Realization of the planning horizon for the first evaluation, green indicates that no lice treatment is necessary, while red indicates the opposite.

Table 9.3 presents the results of our stochastic study for each of the smolt type combinations and their respective candidate solutions. When calculating the average values for a specific candidate solution over its evaluations, we exclude evaluations that turn out infeasible, and report the number of infeasible evaluations as a seperate data point. This is done as the company would in reality manage to solve these infeasibility issues while incurring some minor penalties. Additionally, setting infeasible evaluations' values to zero or $-\infty$ and incorporating them in the candidate solution's average would heavily reduce the interpretability of said average. All of the feasibility issues we encounter are caused by activity constraints requiring a deployment, but the company cannot deploy biomass without breaching the company wide MAB. Moreover, the gap between the best and worst outcome for both the objective values and harvest volumes for all candidate solutions is less than 5 % in all instances. Indicating that the choice of candidate solution does not lead to any catastrophic results, given that a feasible solution is found.

To start off we note that the objective values and harvest volumes have decreased compared to the deterministic no-lice study, which is expected as lice treatments reduce growth rates and increase mortality rates. If we compare the average objective value for the different smolt combinations with the corresponding objective value provided by the RHH relax 2y solution in the no-lice study, we find decreases between 6.4 and 7.2 %. Moreover, we see that the differences between the average objective values and harvest volumes have increased compared to our no-lice study. This is expected as the use of out of sample sampling might require exploitation of the harvest

Table 9.3: Results of the stochastic study. Gap measurements given in parenthesis.

Smolt Types	Cand. Sol.	Avg. Objective Value	Avg. Harvest Volumes	# Infeasible	Avg. HOG Harv. Weight	SRCL	LRCL	Avg. Solution Time [s]
Reg.	1	82,609.13 (2.86 %)	82,618.46 (2.85 %)	0	4,027.27	17.09	18.50	4,325.27 (49.08 %)
Reg.	2	82,562.61 (2.55 %)	82,595.58 (2.63 %)	1	3,948.72	16.99	18.40	4,149.00 (55.20 %)
Reg.	3	82,162.06 (3.49 %)	82,201.67 (3.49 %)	0	3,934.12	17.14	18.43	4,078.87 (43.34 %)
Reg.	All	82,441.92 (3.68 %)	82,469.09 (3.76 %)	1	3,970.52	17.08	18.44	4,185.18 (59.25 %)
Reg. & Mono	1	89,576.80 (2.44 %)	89,584.67 (2.47 %)	0	4,071.29	16.22	18.00	14,553.00 (46.96 %)
Reg. & Mono	2	90,241.45 (2.68 %)	90,266.37 (2.63 %)	0	4,099.53	16.28	17.79	15,992.00 (31.54 %)
Reg. & Mono	3	89,440.69 (2.38 %)	89,445.83 (2.38 %)	0	4,076.53	16.42	17.83	14,395.07 (53.33 %)
Reg. & Mono	All	89,752.98 (3.72 %)	89,765.62 (3.72 %)	0	4,082.45	16.31	17.87	14,980.02 (55.19 %)
All	1	89,760.98 (2.89 %)	89,766.33 (2.84 %)	1	4,076.22	15.97	17.63	34,563.86 (34.34 %)
All	2	90,843.79 (4.71 %)	90,844.38 (4.71 %)	0	4,090.22	16.05	17.61	34,812.67 (27.66 %)
All	3	89,894.52 (2.31 %)	89,912.98 (2.31 %)	1	4,051.16	16.05	17.87	35,494.64 (24.40 %)
All	All	90,182.18 (4.71 %)	90,190.14 (4.71 %)	2	4,072.95	16.02	17.70	34,953.70 (34.34 %)

volume relaxations to achieve feasibility.

To sanity check our results we compare the results of the regular smolt type only combination with historic production data. Eidsfjord Sjøfarm has historically produced around 15,000 HOG tonnes annually, giving $\frac{15{,}000}{0.84} \cdot 5 = 89,285.71$ tonnes of sea weight harvest volumes over five years. Our results show the regular smolt type combination's harvest volumes to be on average 7.6 % lower than the historical harvest volume benchmark. We are contempt with this discrepancy, as the historical benchmark is an approximate number, and fluctuations in production volumes are to be expected over a five year period. We note that the harvest volumes we obtained in the deterministic no-lice study are within a 1.5 % range of the historical benchmark.

As expected, when we increase the amount of available smolt types, both the average objective value and harvest volumes increase. However, as with the no-lice study most, if not all, of the increase can be achieved by introducing mono-gender smolt, and specifically mono-male smolt. When comparing the average objective value over the candidate solutions within the all smolt type combination with the average over the regular smolt type only combination, we see an increase of 9.2 %, which is within 0.3 percentage points of the increase seen in the no-lice study using RRH relax 2y. Indicating that the value of having several smolt types available is fairly similar in the deterministic no-lice and stochastic setting.

We observe that the average harvest weight does not vary much between

Figure 9.2: Breakdown of deployment volumes per smolt type for all realizations of the three different candidate solutions of the all smolt type combination.

the different smolt combinations, but tends to be somewhat higher when multiple smolt types are available. The largest difference in average harvest weight is between the combination with regular and mono-gender smolt and the regular smolt only combination, with an increase of a mere 2.8 % or 112 grams. However, the rearing cycle lengths are clearly reduced when we introduce the novel smolt types. The average short rearing cycle length (SRCL) is reduced by more than a month when all smolt types are available compared to when only the regular smolt type is available, while the long rearing cycle length (LRCL) is reduced by 0.74 months. This indicates that novel smolt types increase production volumes by reducing the rearing cycle lengths, while maintaining similar harvest weights compared to the case where only the regular smolt type is available.

In general we find that the mono-male smolt is the most utilized smolt type when available, with cooled eggs a distant second when available. Figure 9.2 gives the deployment volumes for each smolt type for all evaluations of each

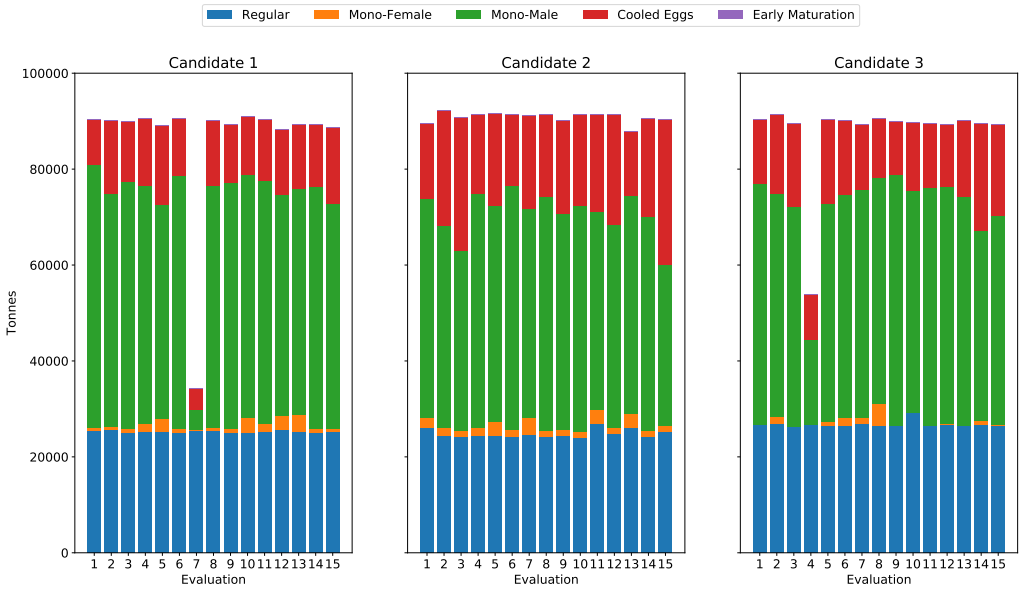

Figure 9.3: Breakdown of harvest volumes per smolt type for all evaluations of the three candidate solutions for the all smolt type combination.

candidate solution of the all smolt type combination. We have included the infeasible evaluations, as they cut of after their infeasibility becomes apparent their values are severely lowered compared to a feasible evaluation. Interestingly, candidate solution 1 and 3 are infeasible for two seperate evaluation runs, meaning that there is not one specifically difficult evaluation run that causes infeasibility. Figure 9.3 provides a similar breakdown for harvest volumes. Note that the regular smolt type's harvest volumes mostly stem from the initial conditions, as only small volumes of regular smolt are deployed during the planning horizon. Moreover, the early maturation smolt is not used at all, and the mono-female smolt is barely used.

Figure 9.4 shows the breakdown of employed biomass for each location and smolt type over the planning horizon for an evaluation of the first candidate solution with all smolt types available. The patterns visible in this specific evaluation are similar to other evaluations of both the same candidate solution and the two other candidate solutions. An interesting aspect to note is the relatively low utilization of the locations Hagebergan, Haukøya Ø, and

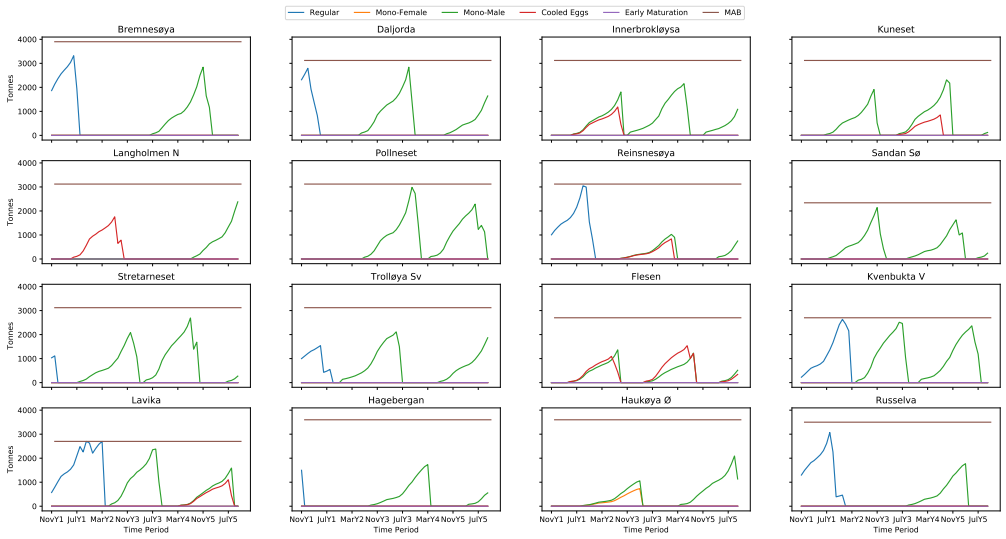

Figure 9.4: Breakdown of employed biomass over time per location and smolt type for an evaluation of the first candidate solution with all smolt types available.

Russelva, which are situated in the coldest and northern most region, which also is the region least affected by lice outbreaks. Showing that the model prefers to deploy smolt at locations with higher temperatures, even though they are more plagued by lice. This indicates that higher growth rates outweighs the higher risk of lice treatments.

9.3.3 Value of Stochastic Planning

In the stochastic study presented above, we use scenario trees to plan for the future. An interesting aspect to examine, is whether this actually adds any value, compared to planning under the assumption that no lice treatments will occur. In this subsection we use only one scenario in our scenario trees with all $\zeta_{it}^s = 0$, except for when we realize the central period where we use the realizations that we sampled in subsection 9.3.2 for the time periods contained in the central period, but we still have $\zeta_{it}^s = 0$ for the forecast period of the scenario. Hence, we choose to run the RHH evaluation scheme 15 times using these realizations of the central periods. However, we choose to not generate a common candidate solution, as the scenario tree used to generate candidate solutions would be equal, i.e. $\zeta_{it}^s = 0$ for all nodes in all of

Table 9.4: Deterministic planning compared to stochastic planning when maximizing harvest volumes.

Smolt Types	Planning	Avg. Objective Value	Avg. Harvest Volumes	# Infeasible	Avg. HOG Harv. Weight	SRCL	LRCL	Avg. Solution Time [s]
Reg.	Stoch.	82,441.92 (3.68 %)	82,469.09 (3.76 %)	1	3,970.52	17.08	18.44	4,185.18 (59.25 %)
Reg.	Det.	81,910.36 (4.07 %)	81,912.82 (4.08 %)	4	3,779.05	14.51	15.34	150.64 (39.90 %)
Reg. & Mono	Stoch.	89,752.98 (3.72 %)	89,765.62 (3.72 %)	0	4,082.45	16.31	17.87	14,980.02 (55.19 %)
Reg. & Mono	Det.	90,231.82 (3.11 %)	90,253.05 (3.11 %)	0	3,855.87	16.03	17.25	589.20 (66.82 %)
All	Stoch.	90,182.18 (4.71 %)	90,190.14 (4.71 %)	2	4,072.95	16.02	17.70	34,953.70 (34.34 %)
All	Det.	91,325.47 (3.17 %)	91,328.26 (3.17 %)	0	3,732.01	15.72	17.39	1,445.20 (41.22 %)

the trees. This can be viewed as generating one candidate solution for each of the 15 evaluation runs, and only evaluating the candidate solution once. To be clear this study differs from the RHH relax 2y applications in the no-lice study, only in that the realization of the central period is now stochastic.

The results of this study are presented and compared to the stochastic study presented in the previous subsection in Table 9.4. As the average objective value of the stochastic solution is never greater than 0.65 % of the deterministic planning solution's average objective value, we conclude there is little value in planning using stochastic programming compared with planning under the assumption that no lice treatments will be necessary. In fact, the objective values even increase on average for the the regular and monogender smolt types combination, and for the all smolt types combination. However, we see that in the case of the regular smolt type combination, the deterministically produced plans turn out infeasible in 4 out 15 evaluations, which is considerably higher than when planning stochastically, where the evaluations turned out infeasible in 1 out 45 evaluations. This indicates that deterministic planning likely plans with much less safety margin with regards to the company MAB.

In general the model seems to plan more aggressively when we use deterministic planning, which is expected as it tries to take advantage of knowing the precise growth and mortality rates a priori. This can be seen through the fact that we now deploy early gender maturation salmon. Early gender maturation smolt has a very short harvest phase consisting of two months where it enjoys its growth boost, however, these two months also have a high risk of lice treatments. Thus, there is a considerable risk of significant biomass losses during these two months. Graphical depictions of the deployment and harvest volumes per smolt types can be found for the all

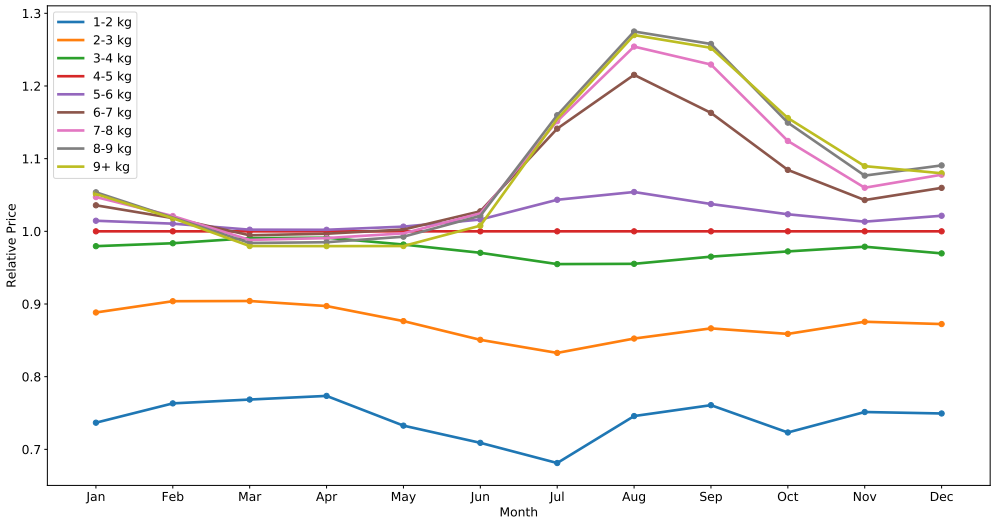

Figure 9.5: Average monthly relative prices computed based on data from NASDAQ Salmon Index for the years 2014-2019.

smolt types combination in Appendix B. Moreover, we see a noticeable drop in rearing cycle lengths and average harvest weights when planning deterministically compared to planning stochastically. For the regular smolt type only combination we see the short rearing cycle length (SRCL) drop by a staggering 2.57 months when the evaluations turn out feasible.

9.4 Maximizing for Relative Price

We now turn to maximizing for relative prices as discussed in section 4.3 and 5.4. We take the monthly relative prices we computed in section 5.4 and average them over the years 2014 – 2019, the results of which are presented in Figure 9.5. The main change in the relative price instances compared to the instances maximizing harvest volumes is that heavier salmon becomes relatively worth more, especially during late summer and fall. In this instance we skip the deterministic no-lice study and head straight to the stochastic study, before finishing off by studying the value of stochastic planning in a relative price setting.

Table 9.5: Results of the stochastic relative price study.

Smolt Types	Cand. Sol.	Avg. Objective Value	Avg. Harvest Volumes	# Infeasible	Avg. HOG Harvest Weight	SRCL	LRCL	Avg. Solution Time [s]
Reg.	1	81,060.11 (3.80 %)	82,310.50 (3.93 %)	1	4,063.67	17.08	18.23	3,868.71 (34.07 %)
Reg.	2	80,832.92 (3.16 %)	82,015.61 (3.35 %)	0	4,118.08	17.12	18.29	4,301.40 (42.70 %)
Reg.	3	80,629.83 (3.00 %)	81,823.90 (3.23 %)	1	4,077.42	17.23	18.44	3,887.64 (64.71 %)
Reg.	All	80,840.76 (3.92 %)	82,049.20 (4.30 %)	2	4,047.61	17.15	18.32	4,025.81 (66.14 %)
Reg. & Mono	1	88,444.66 (2.94 %)	89,256.56 (3.30 %)	0	4,227.85	16.38	17.93	15,503.87 (37.74 %)
Reg, & Mono	2	88,133.43 (2.73 %)	88,787.96 (3.07 %)	0	4,316.79	16.49	17.97	16,139.93 (45.14 %)
Reg. & Mono	3	87,767.47 (2.30 %)	88,376.33 (2.20 %)	0	4,253.26	16.37	17.89	15,382.27 (37.42 %)
Reg. & Mono	All	88,115.18 (3.40 %)	88,806.95 (4.06 %)	0	4,265.97	16.41	17.93	15,675.36 (49.39 %)
All	1	89,443.10 (2.80 %)	90,336.82 (2.41 %)	0	4,131.63	16.19	17.84	38,134.27 (25.13 %)
All	2	88,774.65 (6.20 %)	89,399.31 (6.02 %)	0	4,258.45	16.26	17.83	38,718.67 (29.94 %)
All	3	89,314.57 (2.44 %)	90,117.25 (2.74 %)	0	4,234.12	16.10	17.71	38,514.67 (26.05 %)
All	All	89,177.44 (6.39 %)	89,951.12 (6.52 %)	0	4,208.07	16.18	17.79	38,455.87 (29.94 %)

9.4.1 Stochastic Study Maximizing for Relative Prices

The results of the relative price study are presented in Table 9.5. On a high level the results are very similar to the results we obtained in the stochastic study maximizing harvest volumes. Due to the use of relative prices in the objective function, the objective values of this study should not be directly compared to the objective values obtained in the maximizing harvest volumes study. An interesting thing to note, is that the average harvest weight is not significantly higher than when optimizing for raw harvest volumes, even though heavier salmon is relatively worth more in this study.

In Figure 9.6 we compare the HOG harvest volumes of the relative price study with the maximizing harvest volumes study, broken down per weight class and month when averaged over all evaluations of the all smolt type combination. We observe that there is a small increase in harvest volumes for heavier weight classes during August and September, when heavier salmon fetch relatively higher prices. However, we do not see any significant shift in harvest behaviors. In fact, the two sub-figures are extremely similar, indicating that the increased price achieved by letting the salmon spend more time at sea to grow, seldomly compensates for the increased rearing cycle length, even though we do not account for feeding costs in our case. We also note the general tendency to harvest during months where temperatures and thus growth rates are higher.

Moreover, we observe similar increases in objective values and harvest volumes when we make more smolt types available to the farmer, as we have seen in our previous studies. When breaking down the deployment volumes

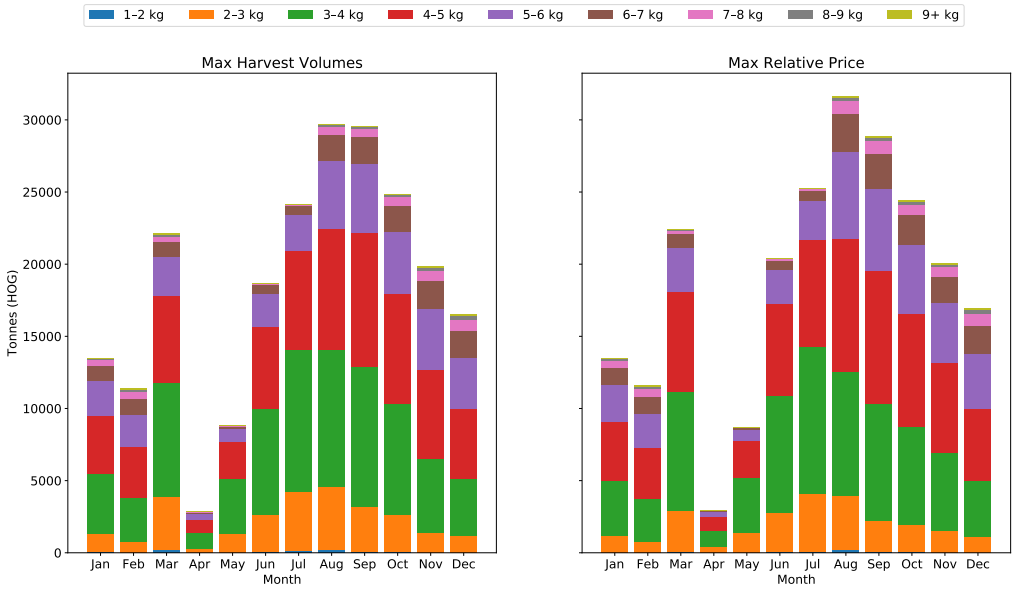

Figure 9.6: Comparison of harvest volumes between maximizing harvest volumes and relative prices, broken down per weight class and month averaged over all evaluations of the all smolt type combination.

per smolt type for the combination where all smolt types are available we see a similar picture as when maximizing harvest volumes, as can be seen in Figure 9.7.

9.4.2 Value of Stochastic Planning in a Relative Price Landscape

We now examine the value of stochastic planning in a relative price setting. This study is similar to the one we performed in the maximizing harvest volumes study. However, as we now maximize with respect to relative prices, the timing of harvest volumes with respect to weight class has increased importance. As such, we expect stochastic planning to have greater value in the relative price setting compared to the case were we simply maximize harvest volumes. As Table 9.6 shows, this is indeed the case. The average objective value is greater when using stochastic planning for all smolt combinations. However, the increase in average objective value is modest, as we see an increase of 0.43 % in the regular and mono-gender combination, and 1.3 % in

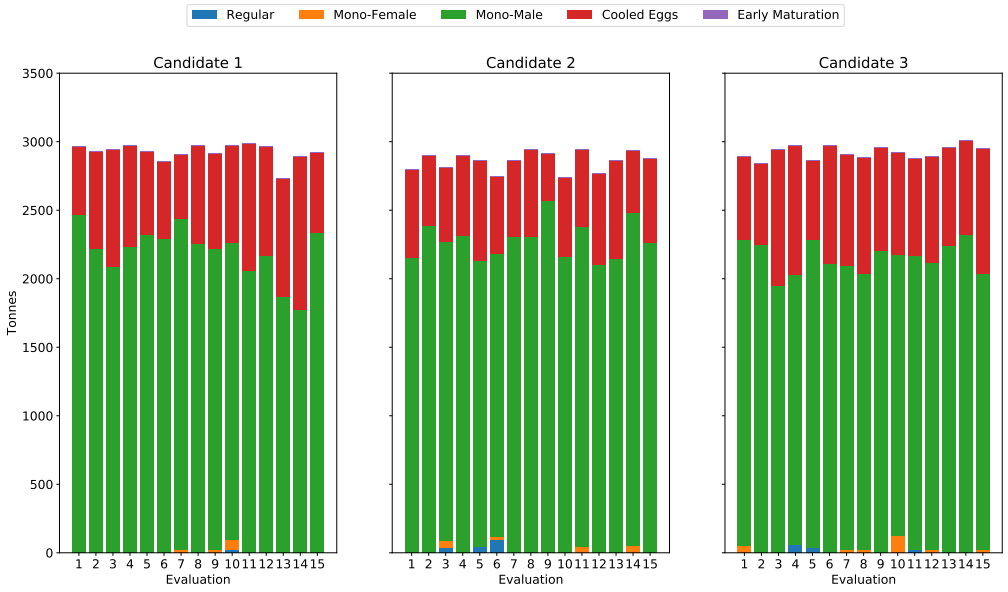

Figure 9.7: Breakdown of deployment volumes per smolt type for all realizations of the three different candidate solutions of the all smolt type combination, when maximizing for relative prices.

the all smolt type combination.

It is clear that using deterministic planning causes a significant increase in infeasibility issues when only the regular smolt type is available. In fact, when using deterministic planning 9 out of 15 evaluations turn out infeasible compared to 2 out of 45 when using stochastic planning. This is also a stark increase compared to the maximizing harvest volumes study, where only 4 out of 15 evaluations turned out infeasible when using deterministic planning. The increase is likely caused by the deterministic plan trying to aggressively allocate harvest volumes of heavier weight classes during August and September. Reducing the available flexibility to change plans when the central period is realized. However, this pattern disappears when we introduce multiple smolt types. We hypothesize that this is attributable to the flexibility obtained through being able to control growth and mortality rates, by picking and mixing between different smolt types differing in growth and mortality rates.

Table 9.6: Comparison of results when performing deterministic and stochastic planning when optimizing for relative prices.

Smolt Types	Planning	Avg. Objective Value	Avg. Harvest Volumes	# Infeasible	Avg. HOG Harv. Weight	SRCL	LRCL	Avg. Solution Time [s]
Reg.	Stoch.	80,840.76 (3.92 %)	82,049.20 (4.30 %)	2	4,077.42	17.15	18.32	4,025.81 (66.14 %)
Reg.	Det.	80,031.08 (2.59 %)	81,437.61 (2.44 %)	9	3,967.27	16.76	17.95	134.50 (33.11 %)
Reg. & Mono	Stoch.	88,115.18 (3.40 %)	88,806.95 (4.06 %)	0	4,265.97	16.41	17.93	15,675.36 (49.39 %)
Reg. & Mono	Det.	87,738.94 (3.79 %)	88,923.16 (3.80 %)	0	4,082.61	16.27	17.47	434.20 (58.08 %)
All	Stoch.	89,177.44 (6.39 %)	89,951.12 (6.52 %)	0	4,208.07	16.18	17.79	38,455.87 (29.94 %)
All	Det.	87,981.59 (3.67 %)	89,362.18 (3.39 %)	0	4,073.66	15.97	17.34	1,481.27 (65.84 %)

Chapter 10

Further Research

In this chapter we outline some potential avenues for further research. The most obvious next step is to apply the model to real world data from seawater trials of the cooled eggs and early maturation smolt type. In the following we discuss avenues that are not mere applications of existing work to new data.

One path of further study is to increase the scope, granularity, detailedness of the model. One natural extension of this thesis is to incorporate different cost aspects in the objective function, such as feeding, deployment and harvesting costs. Another possible extension is to account for a larger part of the value chain, such as the fresh water stage and slaughtering and processing operations. Additionally, increasing the granularity of the model by including the concept of net-pens, and possibly modeling at a weekly or bi-weekly level during the early parts of the planning horizon, could increase the usefulness of the model as a real world decision support tool.

Another dimension that needs improvement is solution methods for the current model. One possibility is to use integer programming methodologies such as column generation and Benders' decomposition, to reduce compute time and thus allow for incorporation of more stages and scenarios in the program. However, we believe it would be interesting to try and utilize machine learning methods such as reinforcement learning to bypass the use of traditional integer programming techniques as branch and cut. One could also develop heuristics to generate feasible solutions as a starting point for a neighborhood search or as a supplement to traditional branch and cut

in an integer program solver, which could reduce the needed solution time drastically.

Chapter 11

Concluding Remarks

In this thesis we formulate a multi-stage stochastic mixed integer linear program for the seawater stage of the Atlantic salmon farming value chain. We develop a biomass development model based on the thermal growth coefficient that account for mortality. Moreover, we developed a scenario generation algorithm for salmon lice treatments, which consider dependencies between different salmon farms. We simplify the multi-stage program into a two-stage stochastic program due to computational constraints. We evaluate the first-stage solutions of this program through a rolling horizon heuristic based simulation framework, and investigate the potential benefits of introducing novel smolt types. Additionally, we examine the value of taking uncertainty into consideration when performing tactical production planning.

We study two different objective functions, one maximizing raw harvest volumes, and the other maximizing harvest volumes weighted by relative price according to weight classes. Our results show that introducing novel smolt types when maximizing for raw harvest volumes can increase the objective values by approximately 9.4 %, while we see an increase of 10.3 % when maximizing with respect to relative prices. The operational plan resulting from maximizing with respect to relative prices shows a small increase in harvest volumes of heavier weight classes during August and September compared to when maximizing for raw harvest volumes. Interestingly, we find that using stochastic programming and accounting for uncertainty in lice treatments during production planning has few benefits with respect to increasing objective values, compared to planning under the assumption that no lice treatments will be necessary. However, using stochastic programming

when planning tends to produce more robust production plans when only the regular smolt type is available, and particularly in the case of maximizing for relative prices with only the regular smolt type available.

Bibliography

M. Aldrin, R.B. Huseby, A. Stien, R.N. Grøntvedt, H. Viljugrein, and P.A. Jansen. A stage-structured bayesian hierarchical model for salmon lice populations at individual salmon farms – estimated from multiple farm data sets. *Ecological Modelling*, 359:333–348, 2017. ISSN 0304-3800. doi: https://doi.org/10.1016/j.ecolmodel.2017.05.019. URL `https://www.sciencedirect.com/science/article/pii/S0304380017300352`.

Magne Aldrin, Bård Storvik, Anja Bråthen Kristoffersen, and Peder Andreas Jansen. Space-time modelling of the spread of salmon lice between and within norwegian marine salmon farms. *PLOS ONE*, 8(5):1–10, 05 2013. doi: 10.1371/journal.pone.0064039. URL `https://doi.org/10.1371/journal.pone.0064039`.

Robert N. Anthony. *Planning and control systems; a framework for analysis [by] Robert N. Anthony*. Division of Research, Graduate School of Business Administration, Harvard University Boston, 1965.

Ragnar Arnason. Optimal feeding schedules and harvesting time in aquaculture. *Marine Resource Economics*, 7(1):15–35, 1992. ISSN 07381360, 23345985. URL `http://www.jstor.org/stable/42629021`.

Frank Asche and Atle G. Guttormsen. Patterns in the relative price for different sizes of farmed fish. *Marine Resource Economics*, 16(3):235–247, 2001. doi: 10.1086/mre.16.3.42629321. URL `https://doi.org/10.1086/mre.16.3.42629321`.

Frank Asche, Helge Bremnes, and Cathy R. Wessells. Product aggregation, market integration, and relationships between prices: An application to world salmon markets. *American Journal of Agricultural Economics*, 81

(3):568–581, 1999. doi: 10.2307/1244016. URL https://onlinelibrary.wiley.com/doi/abs/10.2307/1244016.

Frank Asche, Bård Misund, and Atle Oglend. The spot-forward relationship in the atlantic salmon market. *Aquaculture Economics & Management*, 20: 222–234, 04 2016. doi: 10.1080/13657305.2016.1156192.

Arnfinn Aunsmo, Randi Krontveit, Paul Steinar Valle, and Jon Bohlin. Field validation of growth models used in atlantic salmon farming. *Aquaculture*, 428-429:249–257, 2014. ISSN 0044-8486. doi: https://doi.org/10.1016/j.aquaculture.2014.03.007. URL https://www.sciencedirect.com/science/article/pii/S0044848614001100.

Britt Bang Jensen, Lars Qviller, and Nils Toft. Spatio-temporal variations in mortality during the seawater production phase of atlantic salmon (salmo salar) in norway. *Journal of Fish Diseases*, 43(4):445–457, 2020. doi: 10.1111/jfd.13142. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/jfd.13142.

Barentswatch. Fiskehelse. https://www.barentswatch.no/fiskehelse/, 2021. Accessed: 2020-04-15.

Trond Bjørndal. Optimal harvesting of farmed fish. *Marine Resource Economics*, 5(2):139–159, 1988. doi: 10.1086/mre.5.2.42628926. URL https://doi.org/10.1086/mre.5.2.42628926.

Trond Bjørndal, Daniel E Lane, and Andrés Weintraub. Operational research models and the management of fisheries and aquaculture: A review. *European Journal of Operational Research*, 156(3):533–540, 2004. ISSN 0377-2217. doi: https://doi.org/10.1016/S0377-2217(03)00107-3. URL https://www.sciencedirect.com/science/article/pii/S0377221703001073.

Fernanda Bravo, Guillermo Durán, Abilio Lucena, Javier Marenco, Diego Morán, and Andrés Weintraub. Mathematical models for optimizing production chain planning in salmon farming. *International Transactions in Operational Research*, 20(5):731–766, 2013. doi: https://doi.org/10.1111/itor.12022. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/itor.12022.

J.R. Brett and T.D.D. Groves. 6 - physiological energetics. In W.S. Hoar, D.J. Randall, and J.R. Brett, editors, *Bioenergetics and Growth*, volume 8 of *Fish Physiology*, pages 279-352. Academic Press, 1979. doi: https://doi.org/10.1016/S1546-5098(08)60029-1. URL http://www.sciencedirect.com/science/article/pii/S1546509808600291.

FAO. *The State of World Fisheries and Aquaculture 2020*. FAO, 2020. doi: 10.4060/ca9229en. URL https://doi.org/10.4060/ca9229en.

Martin Faustmann. Calculation of the value which forest land and immature stands possess for forestry. *Allgemeine Forstund Jagdzeitung*, 15:441–455, 1849.

Fiskeridirektoratet. Fiskehelserapporten 2020, 2020. URL https://www.vetinst.no/rapporter-og-publikasjoner/rapporter/2021/fiskehelserapporten-2020.

Odd Inge Forsberg. Empirical investigations on growth of post-smolt atlantic salmon (salmo salar l.) in land-based farms. evidence of a photoperiodic influence. *Aquaculture*, 133(3):235 – 248, 1995. ISSN 0044-8486. doi: https://doi.org/10.1016/0044-8486(95)00029-2. URL http://www.sciencedirect.com/science/article/pii/0044848695000292.

Odd Inge Forsberg. Optimal stocking and harvesting of size-structured farmed fish: A multi-period linear programming approach. *Mathematics and Computers in Simulation*, 42(2):299–305, 1996. ISSN 0378-4754. doi: https://doi.org/10.1016/0378-4754(95)00132-8. URL https://www.sciencedirect.com/science/article/pii/0378475495001328. Mathematical Modelling and Simulation in Agriculture and Bio-Industries Proceedings of the 1st IMACS-IFAC Symposium Msu2SABI.

Odd Inge Forsberg. Optimal harvesting of farmed atlantic salmon at two cohort management strategies and different harvest operation restrictions. *Aquaculture Economics & Management*, 3(2):143–158, 1999. doi: 10.1080/13657309909380241. URL https://www.tandfonline.com/doi/abs/10.1080/13657309909380241.

Odd Inge Forsberg and Atle G. Guttormsen. The value of information in salmon farming. harvesting the right fish at the right time. *Aquaculture*

Economics & Management, 10(3):183–200, 2006. doi: 10.1080/13657300600985215. URL https://doi.org/10.1080/13657300600985215.

Stephen J. Gange. Generating multivariate categorical variates using the iterative proportional fitting algorithm. *The American Statistician*, 49(2): 134–138, 1995. ISSN 00031305. URL http://www.jstor.org/stable/2684626.

Atle G. Guttormsen. Faustmann in the sea: Optimal rotation in aquaculture. *Marine Resource Economics*, 23(4):401–410, 2008. doi: 10.1086/mre.23.4.42629671. URL https://doi.org/10.1086/mre.23.4.42629671.

Terry Heaps. The optimal feeding of farmed fish. *Marine Resource Economics*, 8(2):89–99, 1993. doi: 10.1086/mre.8.2.42629053. URL https://doi.org/10.1086/mre.8.2.42629053.

Terry Heaps. Density dependent growth and the culling of farmed fish. *Marine Resource Economics*, 10(3):285–298, 1995. doi: 10.1086/mre.10.3.42629592. URL https://doi.org/10.1086/mre.10.3.42629592.

Martin B. Hæreid, Peter Schütz, and Asgeir Tomasgard. *A Stochastic Programming Model for Optimizing the Production of Farmed Atlantic Salmon*, pages 289-311. 2013. doi: 10.1142/9789814407519_0011. URL https://www.worldscientific.com/doi/abs/10.1142/9789814407519_0011.

Martin Bergan Hæreid. Allocating sales in the farming of atlantic salmon: Maximizing profits under uncertainty. Master's thesis, Norwegian University of Science and Technology, 2011.

Marianne Iversen, Anne Ingeborg Myhr, and Anna Wargelius. Approaches for delaying sexual maturation in salmon and their possible ecological and ethical implications. *Journal of Applied Aquaculture*, 28(4):330–369, 2016. doi: 10.1080/10454438.2016.1212756. URL https://doi.org/10.1080/10454438.2016.1212756.

George K. Iwama and Arthur F. Tautz. A simple growth model for salmonids in hatcheries. *Canadian Journal of Fisheries and Aquatic Sciences*, 38(6): 649–656, 1981. doi: 10.1139/f81-087. URL https://doi.org/10.1139/f81-087.

Malcolm Jobling. The thermal growth coefficient (tgc) model of fish growth: a cautionary note. *Aquaculture Research*, 34(7):581–584, 2003. doi: https://doi.org/10.1046/j.1365-2109.2003.00859.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1046/j.1365-2109.2003.00859.x.

Tróndur J. Kragesteen, Knud Simonsen, André W. Visser, and Ken H. Andersen. Optimal salmon lice treatment threshold and tragedy of the commons in salmon farm networks. *Aquaculture*, 512:734329, 2019. ISSN 0044-8486. doi: https://doi.org/10.1016/j.aquaculture.2019.734329. URL https://www.sciencedirect.com/science/article/pii/S004484861930609X.

Anja B. Kristoffersen, Daniel Jimenez, Hildegunn Viljugrein, Randi Grøntvedt, Audun Stien, and Peder A. Jansen. Large scale modelling of salmon lice (lepeophtheirus salmonis) infection pressure based on lice monitoring data from norwegian salmonid farms. *Epidemics*, 9:31–39, 2014. ISSN 1755-4365. doi: https://doi.org/10.1016/j.epidem.2014.09.007. URL https://www.sciencedirect.com/science/article/pii/S1755436514000565.

Mowi ASA. Salmon farming industry handbook 2020, 2020. URL https://corpsite.azureedge.net/corpsite/wp-content/uploads/2020/06/Mowi-Salmon-Farming-Industry-Handbook-2020.pdf.

Nærings- og fiskeridepartementet. Forskrift om drift av akvakulturanlegg (akvakulturdriftsforskriften), 2008. URL https://lovdata.no/pro/SF/forskrift/2008-06-17-822.

Nærings- og fiskeridepartementet. Forskrift om bekjempelse av lakselus i akvakulturanlegg, 2012. URL https://lovdata.no/dokument/SF/forskrift/2012-12-05-1140.

Nærings- og fiskeridepartementet. Forskrift om produksjonsområder for akvakultur av matfisk i sjø av laks, ørret og regnbueørret (produksjonsområdeforskriften), 2017. URL https://lovdata.no/pro/SF/forskrift/2017-01-16-61.

Martin Næss and Fredrik Selnes Patricksson. Production planning for atlantic salmon under uncertainty with impact of extensive site school. Master's thesis, Norwegian University of Science and Technology, 2019.

Atle Oglend and Vesa-Heikki Soini. Implications of entry restrictions to address externalities in aquaculture: The case of salmon aquaculture. *Environmental and Resource Economics*, 77(4):673–694, October 2020. doi: 10.1007/s10640-020-00514-0. URL https://doi.org/10.1007/s10640-020-00514-0.

Kathy Overton, Tim Dempster, Frode Oppedal, Tore S. Kristiansen, Kristine Gismervik, and Lars H. Stien. Salmon lice treatments and salmon mortality in norwegian aquaculture: a review. *Reviews in Aquaculture*, 11(4):1398–1417, 2019. doi: https://doi.org/10.1111/raq.12299. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/raq.12299.

Sean Pascoe, Premachandra Wattage, and Debananda Naik. Optimal harvesting strategies: Practice versus theory. *Aquaculture Economics & Management*, 6(5-6):295–308, 2002. doi: 10.1080/13657300209380320. URL https://doi.org/10.1080/13657300209380320.

N.K.G. Salama, A.G. Murray, A.J. Christie, and I.S. Wallace. Using fish mortality data to assess reporting thresholds as a tool for detection of potential disease concerns in the scottish farmed salmon industry. *Aquaculture*, 450:283 – 288, 2016. ISSN 0044-8486. doi: https://doi.org/10.1016/j.aquaculture.2015.07.023. URL http://www.sciencedirect.com/science/article/pii/S004484861530106X.

SinkabergHansen AS. Følg og lær laks fra start til mål!, 2021. URL https://sinkaberghansen.no/folg-og-laer-laks-fra-start-til-mal/.

Statistisk sentralbyrå. Eksport av Laks. https://www.ssb.no/laks, 2021. (In Norwegian) Accessed: 2021-03-07.

Helgi Thorarensen and Anthony P. Farrell. The biological requirements for post-smolt atlantic salmon in closed-containment systems. *Aquaculture*, 312(1):1–14, 2011. ISSN 0044-8486. doi: https://doi.org/10.1016/j.aquaculture.2010.11.043. URL https://www.sciencedirect.com/science/article/pii/S0044848610008161.

Siri Vike. *Infectious salmon anaemia in Atlantic salmon, Salmo salar L. in Chile – transmission routes and prevention.* PhD thesis, University of Bergen, 2014.

Worldometers.info. World's population. https://www.worldometers.info/world-population/#table-forecast, 2021. Accessed: 2021-06-02.

Appendices

Appendix A

Distance Between Locations

Table A.1: Distance between locations in kilometers in the production system used in our case studies.

Location / Location	Bremnesøya	Daljorda	Innerbrokløysa	Kuneset	Langholmen N	Pollneset	Reinsnesøya	Sandan Sø	Stretarneset	Trolløya Sv	Flesen	Kvenbukta V	Lavika	Hagebergan	Haukøya Ø	Russelva
Bremnesøya	0.00															
Daljorda	59.68	0.00														
Innerbrokløysa	21.75	42.87	0.00													
Kuneset	62.93	3.68	46.53	0.00												
Langholmen N	55.81	3.93	39.42	7.25	0.00											
Pollneset	65.07	6.33	49.41	3.14	9.50	0.00										
Reinsnesøya	5.94	53.15	16.08	57.13	49.85	59.92	0.00									
Sandan Sø,	8.10	66.92	29.77	70.85	63.62	73.63	13.98	0.00								
Stretarneset	58.46	2.67	42.26	4.97	2.86	7.03	52.66	66.76	0.00							
Trolløya Sv	61.30	2.33	45.14	1.38	5.87	4.05	55.56	69.65	3.61	0.00						
Flesen	81.82	140.53	102.70	143.47	136.39	145.25	86.49	82.28	138.71	141.56	0.00					
Kvenbukta V	84.57	145.55	104.61	145.61	138.73	148.61	88.92	85.42	141.73	144.89	2.94	0.00				
Lavika	85.07	142.05	104.93	145.92	139.12	148.70	89.02	85.37	141.73	144.59	2.78	2.75	0.00			
Hagebergan	283.28	351.05	303.54	344.19	340.11	346.07	287.26	283.66	340.30	343.37	212.12	217.11	214.47	0.00		
Haukøya Ø	295.46	363.18	315.54	355.89	348.93	359.93	300.13	296.35	352.48	355.43	224.44	229.23	227.41	12.80	0.00	
Russelva	273.15	336.70	292.88	333.92	325.83	335.65	276.25	272.73	329.02	331.84	203.64	205.54	203.75	13.99	26.50	0.00

Appendix B

Deterministic Planning Maximizing Harvest Volumes Deployment and Harvest Volumes

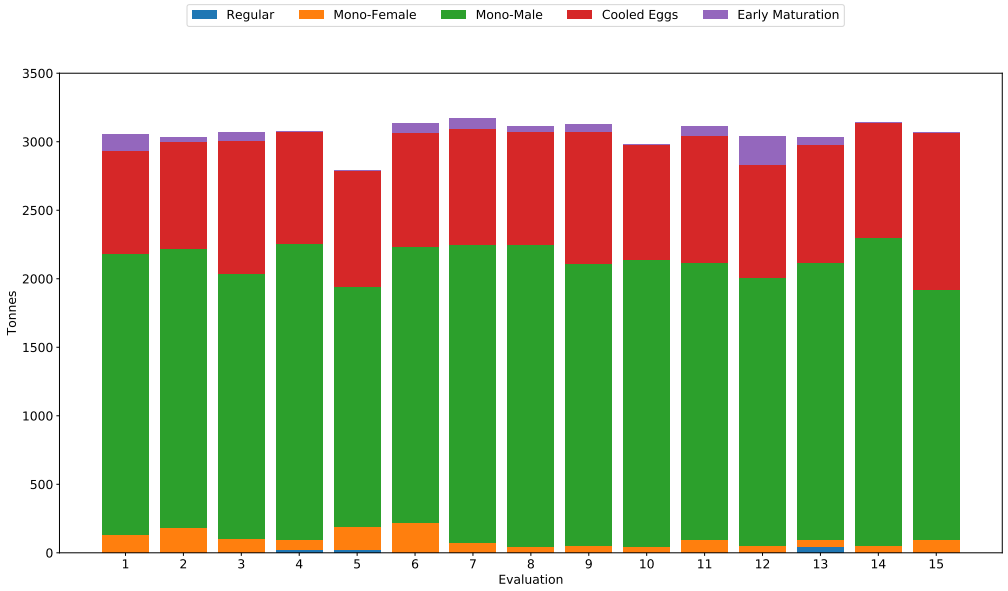

Figure B.1: Breakdown of deployment volumes per smolt type for all evaluation runs of deterministic planning when maximizing harvest volumes with the all smolt type combination.

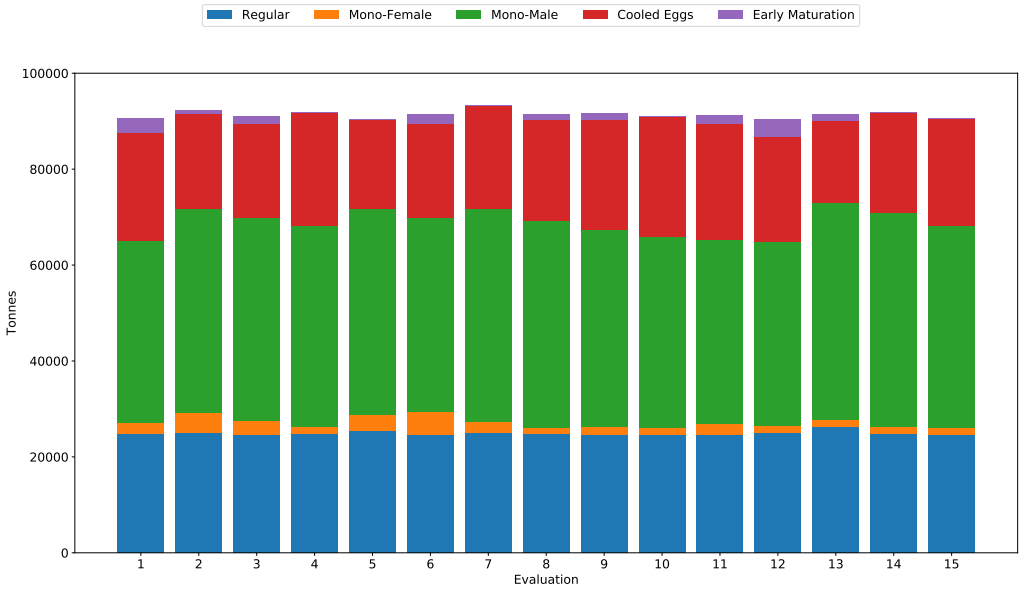

Figure B.2: Breakdown of harvest volumes per smolt type for all evaluation runs of deterministic planning when maximizing harvest volumes with the all smolt type combination.

Randulf Martin Aasen

Tactical Production Planing for Atlantic Salmon Farming Under Uncertainty of Salmon Lice

NTNU
Kunnskap for en bedre verden