Norwegian School of Economics
Bergen, Autumn 2020

# Salmon Price Forecasting

*A comparison of univariate and multivariate forecasting methods*

**Thanesh Tharmarajah & Benjamin Gjesdal**

**Supervisor: Jonas Andersson**

Master thesis, Economics and Business Administration
Major: Business Analytics

NORWEGIAN SCHOOL OF ECONOMICS

# Acknowledgements

This thesis was written as part of a major in Business Analytics at the Norwegian School of Economics (NHH). We have had the privilege of researching an area of great interest and exceptional importance for the Norwegian economy.

The process of writing this thesis for the past semester has taught us how to apply the knowledge gathered over the course of our Master's. It was both challenging and immensely rewarding using our knowledge for an important topic. Although COVID-19 posed certain hardships, it was enlightening to see how all who helped us along the way were able to adapt to support us.

We would like to extend our utmost appreciation and gratitude to our supervisor, Professor Jonas Andersson. He gave exceptional advice and was always quick to help out whenever needed. Lastly, we would like to thank all our family and friends who read through and critiqued our paper countless times and also Knut Henrik Rolland from Kontali Analyse AS for providing valuable insight and data not accessible publicly.

Norwegian School of Economics

Bergen, December 2020

Thanesh Tharmarajah

Benjamin Gjesdal

# Abstract

The salmon industry is becoming an intrinsic part of the Norwegian economy. It is a commercial activity revolving mostly around a single homogenous product. Consequently, salmon farmers and other participants along the value chain can gain substantial insight into how to conduct their business by understanding future spot price movements, primarily since salmon exhibits considerable price volatility. Therefore, it is of great interest to investigate the extent to which time series forecasting can support short- and long-term strategic planning 12 months ahead. Previous research, such as A.G Guttormsen (1999), has shown promising results from applying well known univariate methods. However, most of the studies are outdated, given market changes. Subsequently, this study will focus on partly proven univariate forecasting methods and two multivariate methods regarding Atlantic salmon price forecasting compared to each other and simple benchmarks. The univariate methods are ARIMA and ETS, while the regression methods applied are GAM and LASSO. We chose GAM and LASSO to allow for non-parametric and parametric fit, respectively. The univariate models utilized the spot price of Atlantic salmon, while the multivariate models are supplemented with 20 variables. Each method's accuracy is assessed using mean absolute error and root mean square error for more straightforward interpretability. Results show that univariate ARIMA and benchmark naïve with an STL decomposition outperform GAM and LASSO, suggesting simpler models are perhaps preferable. GAM is superior among the multivariate methods, which can possibly be attributed to it allowing for non-linear relationships. Despite the poor performance, the multivariate models indicate the importance of several variables. Although the models do not provide satisfactory results, it unfolds the possibility of further research using other regression approaches on Atlantic salmon spot price forecasting.

# Contents

# Tables

# Figures

# 1 Introduction

## 1.1 Motivation

Salmon farming has a large impact on Norway and its economy as it is one of the largest industries in the country, generating a significant amount of export value and workplaces domestically. Internationally, the salmon farming industry is important to overcome global challenges. Overpopulation has become a real issue in the world today; the population is growing by 1.05% yearly, which accounts for around 82 million annual increase (Worldometer, 2020). As a result, demand for food is estimated to increase by 50 %, and the demand for animal-based food is expected to increase by 70% by 2050 (Global Salmon Initiative, 2017). This implies that the world's resources would become scarcer and put more pressure on the global food system. The pressure increases on the already overexploited wild fish reserves as well and farmed salmon offers a solution to this problem. To overcome the challenge of a rising population and limited resources, more attention needs to be given to sustainable food production. Aquaculture, and salmon farming especially, is one of the most sustainable food production systems available today, with a carbon footprint per edible kg equal to less than 10 % of the carbon footprint from equivalent amounts of beef (International Salmon Farmers Association, 2018). The salmon industry is one of the most effective food production systems globally. It is estimated to be six times more efficient than beef, four times more efficient than pork, and three times more effective than poultry when measured in edible meat per 100 kg (Solstad, n.d.). In addition to increased pressure on global food systems, poor protein sources and the use of processed meat have increased, which have led to a higher risk of health problems (Cancer Council, 2018). Therefore, it is not enough to only increase the production of protein sources but increase the production of healthy protein sources. The nutritional benefits of salmon, such as the amount of protein, omega-3, and energy, are higher in salmon than the land-based protein sources such as beef or pork (Solstad, n.d.). The sustainable and effective production of nutritious salmon has made the industry an essential part of overcoming the global issues of overpopulation and malnutrition, which states the importance of developing this industry.

Salmon prices have seen increased fluctuations in recent years, impacting the risk management of producers and other entities along with the entire value chain. As the aquaculture industry is an international business industry, many factors affect the salmon price movement, such as currency fluctuations, sea temperature, and sea lice occurrence. Therefore, it is crucial to understand the undercurrents that impact the price and how it will move in both magnitude and direction. Creating a solid forecast model can provide many advantages for the different participants in the value chain. For instance, from late 2016 to early 2017, the salmon price increased by 20-25 NOK/kg over several months (see Figure 5.1). A good forecast model could have discovered such change in advance. If a salmon farming company knew about this price increase earlier, they could have kept their salmon in the cages longer. Thus, the salmon would have grown larger and then sold more volume at a higher price. At this time, the biomass of Norway was recorded to be 705 079 tonnes. Assuming a total weight increase over the delayed harvest time is 1 000 tonnes, the total growth of accumulated revenue for all salmon farmers in Norway is calculated to be NOK 14.1 billion (706 079 * 1 000 * 20) if the salmon price per kg increased by NOK 20. This example illustrates the theoretical potential in the market. The utility of a forecasting tool in a market of such magnitude is undoubtedly significant. Thus, the salmon market becomes very interesting to investigate. The question then begets how to predict the price so the different parties can make informed decisions in the long and short-term.

## 1.2 Research Question

Based on the discussion above and the advantages of proper forecasting tools in this industry, we have formulated the following research question:

*Can the implementation of univariate and multivariate time series forecast methods create solid forecasts for the price of salmon 12 months ahead?*

Therefore, this paper sets out to create univariate and multivariate models that can be employed by participants along the supply chain, whether it be producers, processors, or wholesalers. Such models create value in several aspects of the value chain, such as deciding when to harvest, understanding the profitability of a futures contract, or deciding the amount of smolt release for upcoming seasons. A capable model should provide an example of how to mitigate risk and maximize profits for salmon farmers and others along the value chain. In addition, multivariate models should contribute by indicating which underlying explanatory variable has the strongest predictive power.

# 2 Literature Review & Background

## 2.1 Salmon Farming Industry

Salmon farming is one of Norway's largest industries, accounting for NOK72.5 billion in export (Norwegian Seafood Council, 2020). Moreover, 8340 are employed directly through the aquaculture industry, while 15000 are thought to be peripherally employed through businesses involved with aquaculture activities (Directorate of Fisheries, 2020). Salmon farming makes up over 94% of total Norwegian aquaculture exports (Norwegian Seafood Council, 2020). The largest export market by far is the EU, with a 60% share (MOWI, 2020), followed by Asia, Eastern Europe, and the USA. Since almost all salmon is exported creates substantial exchange rate risk for the companies operating in Norway, especially against the EUR. There are many reasons salmon farming has grown exponentially in both volume and value since its inception in the 1980s. For instance, increased demand has come with a strong trend focusing on healthy foods and sustainability. The commodity has also seen an increase in supply because of improved technology. In essence, the salmon farming industry is vital for the Norwegian economy and thousands of jobs. Consequently, it can be highly valuable for the farming companies and others along the value chain to have accurate forecasts of the price of salmon.

## 2.2 Literature Analysis

The following study is divided into two key components. We try to create superior forecasts for the price of salmon and explore the features with the strongest explanatory power and how the industry can benefit from it. Therefore, in the following chapter, previous literature about forecasting price and price volatility of salmon is reviewed, followed by the corresponding factors and how it is believed to pertain to the price of salmon.

### 2.2.1 The literature on Salmon Price Forecasting and Volatility

The literature on salmon prices is often divided into two separate categories. Firstly, we examine studies that research the direct prediction of salmon prices in different markets. In this case, direct entails research solely predicting the price of salmon. By contrast, the other category explores the volatility associated with the price. The literature on direct price forecasts is scarce and aged, with the notable exception of one recent study (Bloznelis, 2017). Conversely, research papers on volatility are comparatively recent and more prevalent throughout the salmon industry timeframe. Overall, research within these two aforementioned domains is limited to a handful of research papers. As a result, this paper seeks to shed light on a mostly unexplored subject. The following sections will elaborate on the existing literature and its relevance for this dissertation.

#### 2.2.1.1 Forecasting the Price of Salmon

As far as we know, there are only a few papers directly forecasting the price of salmon. These studies utilize numerous econometric and forecasting models. In addition, the papers include a wide variety of different features providing great insight into what to focus on and what to disregard when structuring future research.

The first paper dates back to 1989 (Lin, Herrman, Lin, & Mittelhammer, 1989), here time series and econometric approaches are combined to provide the optimal forecast of Atlantic salmon from 1989-1992. Firstly, a simultaneous equation model is created to investigate the features affecting Norwegian Atlantic salmon's supply and demand. More precisely, the econometric model consists of three structural equations, one representing supply and two describing demand in the US and European market, respectively. Therefore, in a later segment, our paper will look closer into the supply and demand-driven forces of Atlantic salmon, such as those used here. Because of the novelty of the industry, a standalone time series forecast was deemed insufficient. However, a monthly time series analysis was performed to acquire the features' future value, except for features with sufficient data. These two approaches were then combined to forecast the price of Norwegian Atlantic salmon. The results were satisfactory in many ways, however, current research should have access to more data, thereby making forecasting multivariate variables to a large extent obsolete. In other words, the paper indicates essential supply and demand variables, however, the methods used are largely antiquated given the data accessible presently.

Vukina and Anderson (1994) forecasts the price for five separate salmonids species found in Tokyo's wholesale market. Subsequently, the prices are predicted and compared using four state-space models by modelling non-stationary time series. The results found were adequate when measured according to MSE and MAPE. Moreover, the results were surprisingly well concerning predicting the correct direction. However, more research is needed to improve the results. Hence, this study was followed by Gu And Anderson (1995).

Unlike the previous paper, Gu & Anderson (1995) combine OLS used to model the seasonality removal with a state-space, time-series forecasting method to predict the price for the US salmon market. Similar to Vukina and Anderson (1994), this study uses wholesale price indexes for five salmonids. Four models are compared for out-of-sample 3-, 6-, and 12-months to examine the performance. The results exemplify how accounting for seasonal factors significantly improves the forecasting model. Consequently, our model will account for seasonality.

In 1999 A.G Guttormsen (1999) published a paper that focused on short-term 4-, 6- and 8-weeks forecasts to mitigate risk in the industry. Unlike the two previous papers, this one employed six relatively simple and known models. These were the Holt-Winters Exponential Smoothing (HW), Auto Regressive Moving Average (ARMA), Classical Additive Composition (CAD), Vector AutoRegression (VAR), and two naïve methods.

The latest study by Bloznelis (2017) argues all the research mentioned above are obsolete from an empirical point of view. This study used 16 different methods to forecast 1-5 weeks Atlantic salmon spot prices. Only five variables are used. Among them, we find the price of futures. Although Bloznelis argues for their inclusion, chapter 3 elaborates on why this variable is contentions and the reason it is included in our analysis. Every method Bloznelis (2017) uses gets the directional movement right over 50% of the time for all forecasting horizons. K-nearest neighbour gives the best prediction one week ahead, vector error correction model using elastic net regulation for 2 and 3 weeks ahead, and futures prices for week 4 and 5. The gains from a simple naïve benchmark are marginal; therefore, future research is encouraged. Overall, many univariate methods and a few multivariate models are used. However, most of the univariate models are obsolete, while the few multivariate models are either obsolete or used to forecast a short timeframe, such as Bloznelis (2017). Therefore, encouraging the investigation into improved multivariate methods that have not been utilized for forecasting the price of salmon.

#### 2.2.1.2 Salmon Price Volatility

There are several research papers written on the volatility of salmon price, the first being Oglend and Sikveland (2008). This study used a generalized autoregressive conditional heteroscedasticity (GARCH) approach in order to test for volatility. The results showed that higher volatility was reflected in periods with high prices, which will be accounted for in our thesis.

Equivalently, Oglend (2013) employs a GARCH model. This paper suggests that a correlation between volatility and price is due to strong supply and demand conditions. Firstly, there is a significant positive relationship between substitute food prices and volatility supported by rigorous empirical investigation. Additionally, max allowed biomass is also attributed to affecting the price given that increased demand will not be met by increased supply because of constraints in available biomass. As a result, supply and demand factors such as biomass and alternative proteins will be discussed in more detail in chapter 3.

#### 2.2.1.3 Key Points From the Review

This study predicts the spot price of salmon over a 12-month horizon and tries to understand the relationship between the features and the dependent variable. Therefore, moving forward, there are certain aspects that should be taken into account:

- Previous literature focused almost exclusively on univariate models. Therefore, this study will utilize two multivariate regression-based approaches for 12-steps ahead forecasting.
- We will also forecast using previously utilized univariate methods with new data and compare it to the regression-based approaches.
- A closer look at features said to affect the price of salmon.
  - Exchange rates (Lin, Herrman, Lin, & Mittelhammer, 1989).
  - Protein substitutes and biomass (Oglend, 2013).
  - Futures contracts (Bloznelis, 2017).
- Account for seasonality (Gu & Anderson, 1995).
- Account for price volatility (Oglend & Sikveland, 2008).

# 3 Variable Selection

A machine learning model needs explanatory variables or lagged dependent variables to predict and understand what affects the dependent variable (Hyndman & Athanasopoulos, 2018). In our case, the dependent variable is the salmon spot price per kg, while the predictors are chosen through extensive research and personal communication with Knut Henrik Rolland from Kontali Analyse AS[1].

*Standing Biomass*

The standing biomass is defined as the weight or mass of live fish (Marine Institute, 2018). This parameter is usually measured as individuals or tonnes and is divided into generations depending on its weight. Biomass provides a transparent insight into the farming volume, and by relating this amount to the salmon farming cycle, it becomes a good indicator of future short-term supply (MOWI, 2020). Because the standing biomass can be leveraged in many ways depending on the circumstances, it can be expected to positively and negatively impact the price. Optimally, the lags would reflect the generational makeup, primarily because in the short term, the standing biomass for the larger generations are measured close to harvesting time and would therefore be similar to harvesting volume (MOWI, 2020). However, accurate data on separate generations was unattainable. Thus, the lags are based on using vaccine sales as a rough approximation of the respective generations (MOWI, 2020), which equals a 4- and 8-month lag.

*Alternative Animal Protein*

The price of alternative animal proteins affects the demand for salmon, given their nature as substitutes, thereby impacting the price of salmon. There are many possible substitutes. However, Oglend (2013) limits it to some of the most prevalent meats such as poultry, bovine, and ovine. Also, trout will also be included, given its intrinsic similarities to salmon and the second most-produced salmonid (MOWI, 2020). A price increase on alternative protein sources should enhance salmon consumption, which consequently increases the salmon price. A principal component analysis with price movements of alternative meats, cereals, oils, and fishmeal found that meats alone account for 89.54 % of the variation in salmon price volatility (Oglend, 2013), further supporting their inclusion in the multivariate models. It is difficult to

[1] For further information on sources, see appendix A.1 Data Sources

determine the lag structure of alternative proteins. However, this study assumes a rapid but not immediate effect and that certain considerations have to be accounted for. Firstly, wholesalers have to modify their prices according to changes in price movements for the different meats. As a result, end-users will need additional time to adapt their purchasing behaviour. Therefore, a moderate lag of 1 month is utilized.

*Smolt release*

A juvenile salmon is referred to as a smolt and are usually around two to three years old when they are at the release stage. The smolt release stage is when the smolt are transferred from freshwater to seawater cages or net pens. The smolt uses between one and a half to two years after it is released into seawater before harvesting (MOWI, 2020). By including the amount and time of smolt release, we can estimate an expected harvesting volume from one and a half to two years from the point of release. The implication being that a smolt release is a good long-term indicator of future supply with predictive power at lags 18 to 24 months and a negative impact on the price.

*Sea temperature*

The sea temperature is an environmental factor that impacts the duration of the salmon production cycle. Higher sea temperature enhances salmon growth, and the production cycle becomes shorter, while lower sea temperature slows down the salmon growth, which implies longer production cycles. Therefore, large deviations are expected to have an impact on harvesting volume and future supply (MOWI, 2020). One study also indicates that if sea temperatures rise above a specific threshold, the salmon will be stressed, slowing down the production cycle again (Falconer, et al., 2020). As a result, the increased temperature to a certain point should correspond to higher supply and, therefore, lower price. The variation in the sea temperature is expected to impact salmon growth regardless of generation, which means that any major temperature fluctuations would lead to a change in harvest volume for several generations. Thus, we expect variation in sea temperature to affect with a 3 to 5 months lag.

*Harvest*

Harvest is one of the key indicators of short-term supply, given the short-term expiration of salmon. However, the harvest sees a large seasonal variation given many previously mentioned features such as temperature and smolt release. In addition, other factors impact the harvest, such as demand during the high holidays, which pushes up the price, incentivizing further harvest. As a result, there is a reciprocal interplay, which means that the different companies react to increased prices by increasing harvest, which provides the market with much supply having an adverse effect on price. Consequently, it is assumed that it negatively affects the price, and it is important to utilize it as one of the explanatory variables. Since harvest has an almost immediate effect on the supply, it is thought that its lag is no more than one month.

*Feed Consumption*

The feed consumption is the total amount of fish feed consumed in a particular time period. It can be used as a measure of future supply as higher feed consumption is related to a higher quantity of produced salmon, thereby negatively impacting the price. Adversely, it could also impact the price positively since feed is the highest cost in production. Additionally, an increasing trend in feed consumption indicates fish growth, in other words, larger salmons and peaks right before harvesting time. We can therefore also expect changes in supply with a 2 - 4 months lag from consumption peaks. Data on feed consumption is scarce; therefore, we use the *feed conversion ratio* (FCR) as a proxy. It provides a rough estimate of future short-term supply.

*Sea lice & Sea Lice Treatments*

Sea lice are aqua parasites that feed on salmon blood, skin, and slime by attaching onto the salmon flesh, and is one of the biggest aqua parasite problems fish farming industries deal with (Bloodworth, Baptie, Preedy, & Best, 2019). Given its importance, it was found necessary to use both the amount of sea lice and the corresponding treatments as variables. If the sea lice are not controlled, it may cause damage and secondary infection to the salmon, which slows down the production cycle and increases mortality (MOWI, 2020). As a result, we expect a lower future supply in the long-term if a disease outbreak occurs, which is expected to affect the future supply with 12 months lag. However, in the short-term, we expect the supply volume to increase due to earlier harvesting to avoid sea lice damage on the salmon. Such premature harvesting would affect the future supply a lot quicker and is expected to have a 3-month lag. As a consequence, it is assumed that in the short-term, prices will fall because of increased supply, while in the long-term, prices will rise given a prolonged negative production effect on the price.

*Fish Loss*

Every month about 10-20% of fish is disposed of (Norwegian Directorate of Fisheries, 2020), highlighting the waste associated with the farming process. This includes salmon dying for varied reasons, fish being discarded in processing, and escaped fish. As a result, this necessitates a variable that takes this into account. Norwegian farming companies have much to gain from understanding how much this factor plays into the overall price, considering it has such a large margin that can be improved. Since it reduces available short and long-term supply, it is expected that it positively correlates with price. This variable is thought to have an almost negligible lag, given that the disposal of fish has an almost immediate impact on the available supply. Therefore, the lag for food waste is set at 1 month.

*Exchange Rates*

The salmon market is an international market where the exchange rates affect both the salmon farming industry's cost and revenue sides. Studies where the exchange rate has been used as a predictor have found an increase in USD against NOK resulted in higher supply to the US (Lin, Herrman, Lin, & Mittelhammer, 1989). A majority of the raw materials needed to produce fish feed is imported from The US and Europe, which implies the cost of fish feed is dependent on the exchange rates between the local currency and USD or EUR (MOWI, 2020). Underscoring this, a large share of all farmed salmon is exported to the European market, meaning the export price is therefore also dependent on exchange rates (MOWI, 2020). The variation in exchange rates is expected to impact the spot price in the medium term, while the cost side will be affected more in the long-term. Hence, it is reasonable to expect a lag of 6 and 12 months.

*Consumption*

Consumption in different regions is thought to be a reliable indicator of demand for salmon. Therefore, the assumption is that an increase in demand would have a corresponding price increase. The data contains information on the EU, USA, Japan, and others as the last category. Including only the EU and USA, one is left with 1.644 million tonnes in 2019 (K. Rolland, personal communication, 12 October 2020) or the equivalent of 65 % of the market, highlighting the importance of demand in these two markets. Most salmon are exported fresh head-on-gutted. As a result, the expectation is a lag of only 1 month for the consumption to impact the salmon price.

*Futures contract*

Futures contracts were first introduced in 2006 by Fish Pool ASA to facilitate risk management in the salmon farming market. The literature on the validity of futures contracts with regard to spot price is somewhat contradictory. Firstly, MOWI states the importance the futures market might have on the spot price (MOWI, 2020). This is amplified by Ankamah-Yeboah et al. (2017), who argues that the futures market is fully developed, therefore having a direct effect on the spot price. Furthermore, Bloznelis (2017) utilized futures prices as a variable in several different methods, demonstrating its explanatory power. In addition, Oglend (2013) argues that futures contracts affect the price of salmon. However, this stands in stark contrast to a different paper by Asche et al. (2016), which contends that the futures market is underdeveloped, hence not suitable in a forecasting model for the spot price. Nevertheless, the documentation weighs heavily towards including futures contracts in our analysis. Futures contracts limit the future available supply of salmon by locking a certain amount to specific buyers over an extended period (Bloznelis, 2017). As a result, it is expected that futures are positively correlated with the spot price. The lag is hard to determine, however, a rough estimate would be 1-4 months since this is how long most of the futures contract extends (Fish Pool b., 2020). Based on the elaboration from chapter 3, we can segregate the variables into two main categories: supply driving variables and demand driving variables.

*Table 3.1: Overview of features.*

| | Variable name | Number of lags | Measurement unit |
|---|---|---|---|
| **Supply** | | | |
| | Smolt Release | 18 - 24 Months | Number of individuals |
| | Standing Biomass | 4 , 8 Months | Tonnes |
| | Sea Temperature | 3 , 6 Months | Celcius degrees |
| | Feed Consumption | 2 - 4 Months | Tonnes |
| | Harvest (Norway) | 1 Month | Individuals |
| | Harvest (Ex. Norway) | 1 Month | Tonnes |
| | Fish Loss (Norway) | 1 Month | Number of individuals |
| | Sea Lice Treatments | 3 , 12 Months | Share of salmon |
| | Sea lice | 3 , 12 Months | Average amount/Salmon |
| **Demand** | | | |
| | Bovine Index | 1 Month | Index |
| | Poultry Price | 1 Month | USD/kg |
| | Lamb price | 1 Month | USD/kg |
| | Trout Price | 1 Month | NOK/Kg |
| | Consumption EU | 1 Month | Tonnes |
| | Consumption USA | 1 Month | Tonnes |
| | Consumption Japan | 1 Month | Tonnes |
| | Consumption Other | 1 Month | Tonnes |
| **Other*** | | | |
| | NOK/EUR Exchange rate | 6 , 12 Months | NOK/EUR |
| | NOK/USD Exchange rate | 6 , 12 Months | NOK/USD |
| | Futures contract | 1 Month | NOK/Kg |

** Other includes factors affecting either both or one of supply and demand*

# 4 Methodology

Forecasting is a powerful statistical tool that can be used to expand the information basis for decisions in short-term and strategic long-term planning. The forecast accuracy is highly dependent on model selection and how suitable it is to the data. The historical data used to create forecasts are typically affected by a pattern such as trend, seasonality, and/or cycles, as explained in previous literature (see section 2.2.1.1). In the following chapter, we will elaborate on the models we chose and many essential prerequisites of forecasting. The univariate methods are ARIMA and ETS, along with two benchmarks, naïve and rwdrift. The multivariate models are GAM and LASSO.

## 4.1 Univariate Methods

Univariate forecasting models, also known as extrapolative forecasting models, are forecasting methods where the forecast model solely relies on historical data of the forecasting variable itself (Glantz & Mun, 2011). The univariate forecasting methods perform accurate forecasts assuming that seasonal- and trend patterns from the historical data will continue into the future. Thus, the historical data for the forecast variable has high predictive power. These methods are an inexpensive and effective way to create a simple and reliable forecast.

### 4.1.1 Basic Forecasting Methods

Within forecasting, there are several simple methods. These methods are often less accurate than more complex forecasting models, but the purpose of these models is to serve as a benchmark rather than a method option. All new methods will be compared against these simple forecasting methods, and if they do not perform better, it is pointless to use a complex forecasting method (Hyndman & Athanasopoulos, 2018). We will start by looking into some simple forecasting methods.

A simple forecasting method is the naïve method. This method sets the last observation as the forecasted value for $h$ time periods and can be expressed as:

$$\hat{y}_{T+h|T} = y_T \tag{4.1}$$

Another simple method is an extension of naïve, called rolling drift method. This method permits the forecast to decrease or increase with time. The change over time equals the mean difference over the time series (Hyndman & Athanasopoulos, 2018). This can be expressed as:

$$\hat{y}_{T+h|T} = y_T + \frac{h}{T+1}\sum_{T=2}^{T}(y_t - y_{t-1}) = y_T \ + h\,(\frac{y_T - y_1}{T-1}) \tag{4.2}$$

**Decomposition**

A decomposition breaks a time series down to different components, such as trend, seasonality, or cycles. There are several different methods to decompose a time series. In this study, we chose to use seasonal and trend decomposition using Loess, also known as STL decomposition. This method is known for estimating non-linear relationships, copes with any type of seasonality, and is robust against outliers (Hyndman & Athanasopoulos, 2018). The STL decomposition is adjusted by two parameters: seasonal window and trend window, and controls the flexibility of the trend and seasonal factors. In R, we use the stlf function to forecast the benchmark models with an STL decomposition. This function decomposes the time series with an STL decomposition, thereafter, forecasts a seasonally adjusted time series before reseasonalizing the forecast. This implies that the models will be forecasted by the components, and by setting the seasonal and trend window to automatic, it chooses the most suitable parameter value.

### 4.1.2 Exponential Smoothing

Exponential smoothing (ETS) methods are forecasting methods with the weighted average of previous observations, where the weight of each observation will exponentially decrease as the observations get older (Hyndman & Athanasopoulos, 2018). One of the main advantages of exponential smoothing forecasting methods is that they allow us to create quick and reliable forecasts, which is important for forecasting models.

There are several exponential smoothing methods. However, we choose to use the Holt-Winters exponential smoothing. This is because it accounts for trend, seasonality, and level, all of which necessary for our time series. This method builds on the simple exponential smoothing and Holt´s linear trend method. In this section, we will look at our chosen exponential smoothing method.

**Holt-Winters´ Seasonal Method**

$$Forecasting\ equation:\ \hat{y}_{T+h|t} = l_t + hb_t + s_{t+h-m(k+1)} \tag{4.3}$$

$$Level\ equation: l_t = \alpha(y_t - s_{t-m}) + (1-\alpha)(l_{t-1} + b_{t-1}) \tag{4.4}$$

$$Trend\ equation: b_t = \beta * (l_t - l_{t-1}) + (1-\beta^*)b_{t-1} \tag{4.5}$$

$$Seasonal\ equation:\ s_t = \gamma(y_t - l_{t-1} - b_{t-1}) + (1-\gamma)s_{t-m} \tag{4.6}$$

Where $l_t$ represents the level of a time series at time *t,* $b_t$ represents the estimated trend slope. The seasonal component is denoted as $s_t$ and where $m$ represents the frequency of seasonality. The forecast equation states that the forecasted value $h$ steps ahead are the sum of the last estimated level, the estimated trend times number of steps ahead forecasted, and seasonal component $s_t$. α is a smoothing parameter with a value between 0 and 1, which represents the weight distributed to an observation. The smoothing parameter, therefore, controls the decreasing rate of weights distributed. The estimated level at time *t* is the weighted average of $y_t$ and the sum of the previous level and trend estimation weighted by the smoothing parameter. Lastly, it is seasonally adjusted by subtracting the seasonal component. The trend is estimated at time *t* by the difference in level between time *t* and *t-1* weighted by the smoothing parameter $\beta^*$ and the weighted average of the previously estimated trend (Hyndman & Athanasopoulos, 2018). The seasonal equation states that the seasonal component is a weighted average of the current seasonal index and the seasonal index from the last seasonal time period, *m* time periods ago. The seasonal component is smoothed by the seasonal smoothing parameter $\gamma$ which takes values between 0 and 1 (Hyndman & Athanasopoulos, 2018).

The exponential smoothing method components can be divided into different categories depending on the characteristics of the data. The seasonal component could be divided into none, additive or multiplicative. The trend component can be divided into none, additive or additive damped, and lastly, the error component could be either additive or multiplicative. For simplicity, we will use a built-in argument in the ets function in R, which chooses the most suitable model automatically.

### 4.1.3 ARIMA Models

ARIMA models bases their judgment on the autocorrelations in the data (Hyndman & Athanasopoulos, 2018). The ARIMA model is often referred to as a good forecast model when the dataset is relatively short, and the model is used to forecast a short-term forecast (Adebiyi, Adewumi, & Ayo, 2014). More complex forecasting methods with higher flexibility often require a large training sample and should be forecasted for extended periods. ARIMA is an acronym for AutoRegressive Integrated Moving Average model, which is a combination of the following models, where integrated is the reverse of differencing.

***Autoregressive***

Autoregressive models forecast a variable based on previous observations of the variable itself rather than based on a set of predictors (Hyndman & Athanasopoulos, 2018). The model can be written as:

$$y_t = c + \phi_1 y_{t-1} + \phi_2 y_{t-2} + \cdots + \phi_p y_{t-p} + \varepsilon_t \tag{4.7}$$

This model is referred to as the AR(p) model, where $p$ represents the order of the autoregressive model and represents white noise. Changes in the parameter $\phi$ lead to different time series patterns, while the difference in variation of the error term only changes the time series scale.

**Moving Average**

Moving average models uses historical forecast errors to create what can be described as a regression-like model (Shumway & Stoffer, 2016). The moving average model weights the current value of the time series with previous forecast errors denoted by $\theta$. This model can be expressed as:

$$y_t = c + \varepsilon_t + \theta_1 \varepsilon_{t-1} + \theta_2 \varepsilon_{t-2} + \cdots + \theta_q \varepsilon_{t-q} \tag{4.8}$$

This model is referred to as MA(q) model, where q is the order of the moving average model and is the weight distribution parameter to the previous forecast errors. Since we do not observe values for this is not a conventional regression. As in AR(p) model, adjustments in the $\theta$ parameter cause different time series patterns, while adjustments in the variance of the error term only affect the scale of the time series (Hyndman & Athanasopoulos, 2018).

**Seasonal ARIMA**

The ARIMA model is capable of handling a variety of seasonal data; in such situations, seasonal differencing is always applied first before eventually applying first differencing if needed. In order to create a seasonal ARIMA model, we need to add a seasonal term to our existing non-seasonal ARIMA model. This seasonal term can be expressed as:

$$ARIMA(p, d, q)(P, D, Q)m \qquad (4.9)$$

Where p represents the order of the autoregressive part, d represents the degree of first differencing involved, and q represents the order of the moving average part (Hyndman & Athanasopoulos, 2018). The capital letters P, D, Q are the seasonal notations for the model, and $m$ is the number of observations per year (Hyndman & Athanasopoulos, 2018). The seasonal ARIMA model can be expressed through backshift notation in the following way:

$$(1 - \phi_1 B)(1 - \Phi_1 B^m)(1 - B)(1 - B^m)y_t = (1 + \theta_1 B)(1 + \Theta_1 B^m)\varepsilon_t \qquad (4.10)$$

The backshift operator *B* shifts the data back one period, which means $By_t = y_{t-1}$, and $B^m y_t = y_{t-m}$. If we look at the seasonal ARIMA model from the left, we start with AR models non-seasonal expression, then the seasonal expression, this is followed up by non-seasonal and seasonal difference, and lastly, the MA models non-seasonal, seasonal expression and the error term at the right side of the equation. Here we see the additional seasonal components are multiplied with the non-seasonal components.

The ARIMA model is determined by a non-seasonal term (p,d,q) and a seasonal term (P,D,Q). The number of *first differencing* and *seasonal differencing*, elaborated later in this chapter, represents the values of d and D, respectively. It now remains to determine the values of p, q, P, and Q. There are two alternatives to determine these values: The first alternative is to use the auto.arima function in R, which automatically selects these values and returns a model. The other alternative is to determine these values manually by using the acf and pacf functions in R, which visualizes an autocorrelation plot and a partial autocorrelation plot.

When values for p, d, and q are defined, we need to estimate the parameters *c,* f, q for each order. The maximum likelihood estimation (MLE) method finds the values for the parameters, which maximizes the probability of replicating historical observations (Hyndman & Athanasopoulos, 2018). This method is also applied in R, where the software uses the logarithm of the probability of the data belongs to the estimated model, which means it maximizes the log-likelihood to estimate the parameters (Hyndman & Athanasopoulos, 2018).

The estimated ARIMA model parameters have some conditions displayed in Table 4.1, and if $p$ or $q$ has a higher order than 2, the restrictions are much more complicated. R takes these restrictions into count when estimating the parameters.

*Table 4.1: ARIMA model parameters´ conditions.*

| p | AR ( p ) condition |
|---|---|
| 1 | $-1 < \phi_1 < 1$ |
| 2 | $-1 < \phi_2 < 1, \phi_1 + \phi_2 < 1, \phi_2 - \phi_1 < 1$ |
| **q** | **MA ( q ) condition** |
| 1 | $-1 < \theta_1 < 1$ |
| 2 | $-1 < \theta_2 < 1, \theta_2 + \theta_1 > -1, \theta_1 - \theta_2 < 1$ |

**Stationarity**

An important part of ARIMA and regression is to have a stationary time series. This means the properties of the time series does not depend on the time of observation (Hyndman & Athanasopoulos, 2018). A non-stationary process could be random walks, trends, or cycles, which affect the data's properties over time. Unit root tests are used to check whether the data is stationary or non-stationary, or it can be checked through visual inspection. As a result, two unit root tests and the autocorrelation function will be used to examine the variables' properties in the dataset. Differencing is then used to solve for non-stationary time series.

**Differencing**

One method to change a time series from non-stationary to stationary is called differencing. Differencing can be applied by computing differences between consecutive observations. This calculation removes levels between observations; hence it reduces or eliminates patterns as seasonality and trend (Hyndman & Athanasopoulos, 2018). This differencing method can be written as follows:

$$\Delta y_t = y_t - y_{t-1} \tag{4.11}$$

Applying differencing to a time series implies that the time series data now only consist of $T - 1$ observations because the first observation does not have any previous observations. If the differencing is successful, we can assume the remaining of the time series is approximately white noise, denoted as $e_t$. In some cases, the time series will still be non-stationary after differencing the time series. On such occasions, it is necessary to apply a second-order differencing, where we apply differencing on the differenced time series $\Delta y_t$. In cases with high seasonality, seasonal differencing is preferable. Here we compute the difference between an observation and the previous observation from the same season rather than consecutive observations. Seasonal differencing can be written as:

$$\Delta y_t = y_t - y_{t-m} \tag{4.12}$$

Where *m* represents the number of seasons, the selection of differencing methods is subjective to some extent. For seasonal data, the seasonal differencing should be applied first and eventually first differencing in addition if the data is still non-stationary after seasonal differencing. But there are more objective ways to decide if differencing is required, such as unit root test.

*Kwiatkowski-Phillips-Schmidt-Shin Test*

A well-known stationarity test is named *Kwiatkowski-Phillips-Schmidt-Shin (KPSS)* test. This test the time series to be stationary as the null hypothesis and look for evidence for the null hypotheses to be false (Kwiakowski, Phillips, Schmidt, & Shin, 1992). In R, we activate the "urca" package, which includes a function named ur.kpss. This function computes the KPSS test on the selected time series and returns critical values for different levels of significance and a test statistic. If the test statistic is greater than the critical value of the chosen significance level, the null hypothesis is rejected; thus, the time series is non-stationary.

*Augmented Dickey-Fuller Test*

The Augmented Dickey-Fuller test (ADF) is a unit root test that tests for stationarity in a regression. The null-hypothesis states that there is a unit root or present in the time series univariate or that the data is non-stationary, while the alternative hypothesis is that the variable in question is stationary or trend-stationary. The null hypothesis is rejected if the ADF statistic is less than the critical value. The test is conducted using the following expression:

$$\Delta y_t = \alpha + \beta t + \gamma y_{t-1} + \delta_1 \Delta y_{t-1} + \ldots + \delta_{p-1} \Delta y_{t-p+1} + \varepsilon_t \tag{4.13}$$

Where $\alpha$ is the constant term, $\beta$ denotes the coefficient for a certain time trend, while p is the amount lags included, which is decided when conducting the test (Dickey & Fuller, 1995).

*Autocorrelation Function*

An autocorrelation function (ACF) displays how a time series correlates with x amount of preceding lags p. A stochastic process shows the ACF at zero. The autocorrelation increases with an increase in correlation between lags. The interpretation of an ACF plot is the following: The correlation corresponding to each lag is considered significant if it is higher than the blue dashed line for positive correlation and conversely for negative correlation. In addition, there is a partial autocorrelation function which gives partial correlation. ACF and PACF can in simple terms be expressed as:

$$ACF = \frac{Cov(y_t, y_{t+k})}{\sqrt{Var(y_t)Var(y_{t+k})}} \quad PACF = \frac{Cov(y_t, y_{t-2}|y_{t-1})}{\sqrt{Var(y_t|y_{t-1})Var(y_{t-2}|y_{t-1})}} \tag{4.14}$$

Where $y_t$ is the time series at time t and $y_{t+k}$ is the time series at lag k. Both plots ACF and PACF are created in R with the acf and pacf functions.

# 4.2 Multivariate Methods

Our study employs two multivariate models, LASSO and GAM. There are several reasons why these two were chosen. Chiefly, GAM har allows us to fit and analyze potentially non-linear relationships, while LASSO is an exclusively linear method, thereby allowing for comparing significantly different methods.

## 4.2.1 Least Absolute Shrinkage and Selection Operator

The Least Absolute Shrinkage and Selection Operator (LASSO) is a regression method that performs regularization and variable selection so as to achieve improved interpretability and prediction accuracy (Tibshirani, 1996). This can be expressed as:

$$\sum_{i=1}^{n}(y_i - \beta_0 - \sum_{j=1}^{p}\beta_j x_{ij})^2 \; + \; \lambda\sum_{j=1}^{p}|\beta_j| = RSS + \lambda\sum_{j=1}^{p}|\beta_j| \tag{4.15}$$

The first part represents standard least-squares linear regression, while the second part is the penalty term performing the regularization.

$$\lambda\sum_{j=1}^{p}|\beta_j| \tag{4.16}$$

Where the lambda is a tuning parameter, while beta is the coefficient for the explanatory variables, the lambda determines fit and variable selection. As a consequence of the absolute value, LASSO returns sparse models, with a subset of the variables. When lambda equals zero, LASSO becomes a simple linear regression, however as lambda approaches infinity, the penalty for each coefficient will increase to the point where all variables equal zero. In other words, it is rooted in the bias-variance trade-off. This entails an increase in lambda reduces the flexibility of the fit, thus increasing bias and decreasing variance (James, Witten, Hastie, & Tibshirani, 2017). The lambda value is optimized through rolling-origin (see section 4.4.3).

LASSO regression has many qualified reasons for why it could be capable of providing valuable forecasting for the spot price of salmon. Firstly, it is one of the most interpretable regression methods, also for those without comprehensive knowledge of the field. Secondly, it can provide an extensive decrease in variance at the expense of a slight increase in bias. Moreover, LASSO increases model interpretability by selecting only variables associated with the response. This is especially important given that there are 20 explanatory variables and augmented with 32 of their lagged versions.

### 4.2.2 Generalized Additive Model

The generalized additive model (GAM) was chosen to be utilized for one of our multivariate forecasting methods. In general terms, GAM is a non-parametric extension of linear regression. This entails modeling a univariate response $Y$ with $k$ amount of predictors (James, Witten, Hastie, & Tibshirani, 2017). More precisely, the model takes the sum of the optimized smooth functions, which can be expressed as:

$$y_i = \beta_0 + \sum_j s_j(x_j i) + \varepsilon_i \tag{4.17}$$

Every time series variable component x is fitted using splines or smooths, which are functions composed of several additional functions called basis functions. When modeling each basis function $b_p$ has corresponding coefficients $\beta_p$. The resulting spline is the weighted sum of the basis function. This can be shown as:

$$s(x) = \sum_{k=1}^{K} \beta_l b_k(x) \tag{4.18}$$

The number of basis functions is hard to determine. However, it should be large enough to allow for enough flexibility in the model. The "mgcv" package in R allows for extensive customization, including K amount of basis function for each feature. If K is too small, then the model would not capture complexity in the data. However, the negative side of too large K is it creates computational difficulties. In order to fit the data optimally without overfitting, one has to maximize the penalized log-likelihood, which can be expressed the following way:

$$L_p(\beta) = L(\beta) - \frac{1}{2}\lambda\beta^T S\beta \tag{4.19}$$

The $L_p(\beta)$ is in the penalized log-likelihood, while $L(\beta)$ is the log-likelihood and $\frac{1}{2}\lambda\beta^T S\beta$ is the penalization. The last part is determined by lambda. The wiggliness is denoted as the integral of the squared second derivative over the whole range of the covariate x. Conveniently this can be written as a function of the coefficients, where $\beta^T$ is a vector with all coefficients of the basis functions in the spline, S is the penalty matrix created so that when multiplied by the coefficients gives the original expression.

$$\beta^T S\beta = W \tag{4.20}$$

The lambda is used to regulate the trade-off in order to discover the spline which maximizes the penalized log-likelihood. For instance, if the lambda is too large, then the line eventually becomes a simple OLS. Conversely, if the lambda is exceptionally small, then the data will allow for an overly complex fit. The purpose of this is to optimize the fit or wiggliness. A more complex model renders larger penalization. In other words, one needs the data to be fit without it becoming overly complex. This is done by specifying the REML argument in the gam function in R. REML is an alternative to using prediction error criterion like generalized cross-validation (GCV) and AIC. This is a Bayesian approach and preferable in our case, especially given that the GCV is not programmed for time series forecasting.

There are multiple reasons GAM is a good choice for predicting the salmon spot price. Firstly, GAM allows for a non-linear fit to each predictor. This is important when the data seems to indicate a non-linear relationship between the response and explanatory variables. However, at the same time, GAM does not restrict linear relationships if this is closer to the actual representation of a predictor. Furthermore, this paper wished to both rely on good predictions so industry players can make improved decisions while at the same time providing a model that has a high degree of interpretability. For GAM, one can break down the contribution of each predictor on the response as a way to better visualize each variable. In essence, GAMs are the consequential middle ground between simple linear regression and a black-box model, such as neural networks.

#### 4.2.2.1 Forward Stepwise Selection

In order to achieve the optimal results from fitting a GAM model, one has to subset the variables which are relevant for the time series prediction of the spot price of salmon. For this task, a forward stepwise selection procedure was employed tailored for GAM. In General terms, forward stepwise selection involves starting with only the constant term and no variables. For its consecutive step, the selection is tested using a specific criterion (see section 4.4.1), then either adding or removing the variable based on whether the overall criterion improves. Alternatively, one could use best subset selection, backward or hybrid approach. Best subset would have been preferable because this method would have checked through every available model. In other words, an exhaustive approach. However, best subset is predicated on having a smaller subset of features. The reason for this is that when the feature space expands beyond 20, it will become computationally infeasible, given that the number of models to fit is $2^p$. The data used in this thesis is larger than 20 features when including all their lags. As a consequence, the forward stepwise selection method was chosen because there is a somewhat exhaustive forward step selection method package in R, which in addition to traditional forward stepwise selection uses techniques for bootstrap resampling to check for if significant effects of non-selected features impact the model (Sestelo, Villanueva, Meira-Machado, & Roca, 2016).

## 4.3 Data Transformation and Processing

Data transformation of historical data is often needed in order to create simpler forecasting. These adjustments try to remove variation and increase the consistency for all the data. These are necessary steps because it most likely results in increased accuracy of the forecasts, according to Hyndman & Athanasopoulos (2018). In the following section, we evaluate the sample size needed and prerequisites before the data is applicable for time series forecasting, the corresponding transformation, and cross-correlation.

**Sample Size**

Sample training data is essential for fitting superior models. Especially when there are many parameters to consider, in general, precision is expected to increase with additional observations (James, Witten, Hastie, & Tibshirani, 2017). Hyndman & Athanasopoulos (2018) assert there is no "magical" minimum number of observations needed, in fact, they explain how observation required is a product of the number of predictors estimated and stochasticity in the sample.

Firstly, most statistical modelling necessitates that there are less predictors than observations, with the exception of LASSO in our case. Furthermore, according to Hyndman & Kostenko (2007), processes with substantial variation require more data than once with less variation. For both ARIMA and ETS models, one needs a minimum amount of data. For instance, with ETS, there are seasonal, trend, and level elements that require initial values. For the seasonal part, one needs eleven parameters for the initial element, while two parameters are connected to start trend and level values. This entails a prerequisite of 17 observations for monthly data (Hyndman & Athanasopoulos, 2018). The rationale is similar for ARIMA, where the minimum number of observations is equal to $P + Q + mD + d + q + p + 1$ according to Hyndman & Kostenko (2007). For multivariate regression models, more observations usually result in a better fit and out-of-sample outcome (James, Witten, Hastie, & Tibshirani, 2017). The extent of how many observations are necessary for multivariate is highly contested. However, a recommended heuristic states that one should have 10-20 observations per parameter (Harrel, Lee, Califf, Pryor, & Rosati, 1994).

**Transformation**

Log transformations are used to reduce the variation across the time series. In other words, if the volatility of the data increases with the level (i.e., time series), it would require a transformation as a prerequisite before forecasting. It can simply be expressed as the log across the entire time series x.

$$log(x_i) \tag{4.21}$$

**Cross-Correlation Function**

Sample cross-correlation function (CCF) is the correlation between two time series univariates. The function is advantageous in discovering lags between the response and predictors, which might be useful in predicting the independent variable (Penn state, 2020). Lags are a central part of defining the time series relationship of the data. As a result, several industry considerations were presented in chapter 3 as reasons for suggested lags in predictors. However, sample cross-correlation function gives mathematical support to the lag structure.

## 4.4 Model & Method Evaluation

Fitting and evaluating forecasts is essential in order to achieve the finest results. Therefore, in the following section, several different avenues for determining metrics used in model (variable) selection, parameter tuning, and method comparison are proposed.

### 4.4.1 Akaike's Information Criterion

There are numerous metrics for determining the optimal variable selection, such as R squared, BIC, AIC, and AICc. The aim of each of these metrics is to achieve the most parsimonious model. However, this thesis will use the one best suited for our data, AICc.

Akaike's Information Criterion tries to find the best model by maximum likelihood. It gives an account for each model and the relative quantity of information lost. A lower value for AIC entails a better model compared to others with regard to fitting the data and avoiding unnecessary complexity. AIC is defined as:

$$AIC = T * log\left(\frac{RSS}{T}\right) + 2(k + 2) \tag{4.22}$$

Where T is the number of observations, k is the number of predictors, and RSS the fit of the model. However, AIC has its limitations. When there is a small sample size, then there is a considerable probability for AIC will be overfitted by selecting an excessive amount of features (McQuarrie & Tsai, 1998). As a result, a slight modification has to be made given a small sample size. A small sample size will, in this regard, follow the heuristic proposed by (Burnham & Anderson, 2002), where n/k < 40, which is equivalent to the parameters exceeding 2-3% of the data. Consequently, this thesis uses AICc. It can be expressed as:

$$AICc = AIC + \frac{2k^2 + 2k}{n - k - 1} \tag{4.23}$$

Where in the number of observations and k is the number of parameters. In essence, AICc is a penalty term for each additional parameter.

### 4.4.2 Performance Metrics

There are many ways to measure forecast performance between different models. Two of the most widely used are mean absolute error (MAE), and root mean squared error (RMSE).

Firstly, MAE explains the average absolute error between the measured and the actual values. This metric is calculated in a time series by taking the absolute difference between the measured response $\hat{y}_{t\ +h|t}$ with the actual observation $y_{t\ +\ h}$, shortened to the error e.

$$MAE = \frac{1}{T}\sum_{t=1}^{T}\left|e_{t+h|t}\right| \tag{4.24}$$

Next is RMSE, which is an extension of the mean square error (MSE). RMSE is a measure of the root of the squared errors. In other words, MSE only measures the difference between the squared observed and the predicted value. To calculate the RMSE, one needs to take the square root. This is done in order to have output equivalent to the data. Consequently, RMSE is simpler to interpret compared to the MSE. In addition, negative and positive values do not cancel each other.

$$RMSE = \sqrt{\frac{\sum_{t=1}^{T} e_{t+h|t}^{2}}{T}} \tag{4.25}$$

The main difference between the two methods is that RMSE weights larger errors higher, while MAE weights all errors equally. As a result, RMSE is more exposed to outliers than MAE. Both models will be utilized in chapter 6.

### 4.4.3 Nested Cross-Validation

Traditional cross-validation techniques are unreasonable when dealing with time-series forecasting. This is a result of the sequential nature of time series; each observation is associated with the preceding and succeeding month. However, to avoid this, one can implement nested cross-validation.

Within nested cross-validation, there is an approach called rolling-origin. This method divides the data into an outer- and inner-loop. Firstly, a test set is left out. Since this set is used to validate the performance of the fit, it should correspond to the forecast horizon of 12 months (Hyndman & Athanasopoulos, 2018). This is a consequence of the need for a valid estimate of the performance on data not used in the fitting. The remaining data also called the inner-loop, is used to train the model. This training subset is then temporarily split into a training- and validation-set of a given size. The training set is usually 50% or more of the training subset observations, while the validation set is equivalent to the forecast horizon. Furthermore, to make the best use of the data, each fit is moved chronologically forward a single observation as illustrated in  Figure 4.1.

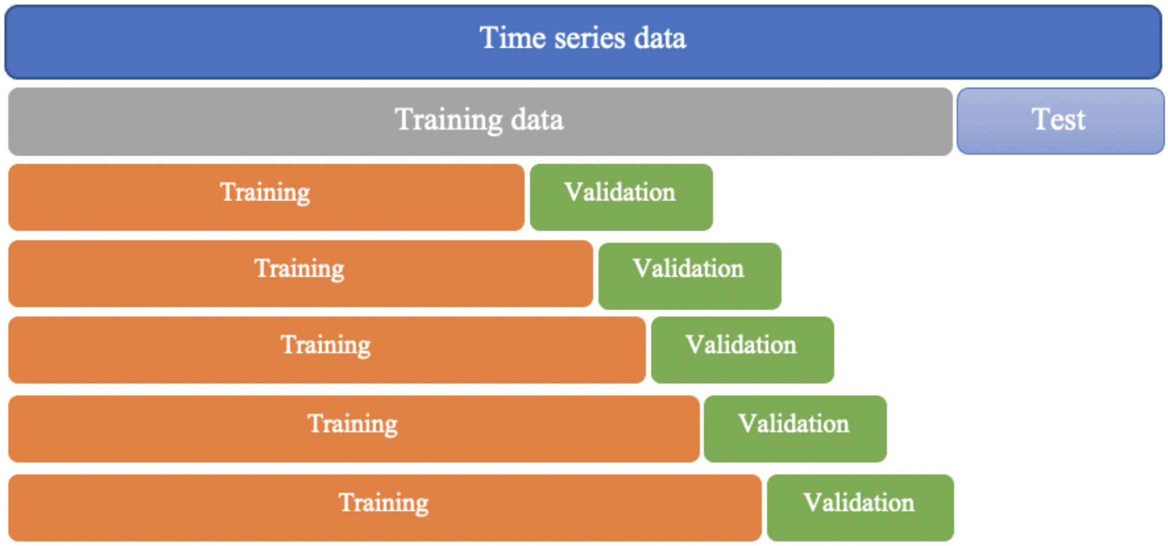


*Figure 4.1: Illustration of rolling-origin nested cross-validation.*

The grey line represents the inner loop data, while the additional light-blue rectangle shows the test hold-out set. In this thesis, cross-validation will be used to fit the models and check for accuracy. Fitting is done by averaging the parameters or coefficients from every fit for a 12-step ahead forecast horizon. When checking the accuracy, every fit is used to compute a forecast. Each forecast is measured using a metric such as MAE and RMSE, which is consequently averaged out for each forecast. This approach is applicable both when dealing with the univariate models and multivariate models.

# 5 Data Analysis and Processing

In order to model the data, several initial steps had to be taken. Firstly, the data was cleaned and fitted to the same frequency. Thereafter, descriptive statistics and statistical properties tests were used to check for stationarity, autocorrelation, cross-correlation, and other factors that affect the data. The following chapter explores these steps in detail.

## 5.1 Fish Pool Index

For the independent variable, we use the Fish Pool Index (FPI) as a reference for the spot price of Atlantic salmon, denoted in NOK/kg. This index is based on a weekly weighted average of the different three weight classes of head-on-gutted (HOG) salmon. These weight classes are divided into 3-4kg, 4-5kg, and 5-6kg, which are weighted 30%, 40%, and 30%, respectively (Fish Pool a., 2020). Subsequently, monthly prices are calculated by averaging the weekly prices, thus what will be referenced hereafter. The index was chosen because salmon contracts in Norway are almost exclusively sold through this spot price reference by exporters. The index has throughout its existence weighted varying underlying reference prices, the most prevalent being NASDAQ salmon index weighted at around 85-90% and SSB export prices at 5-10%. All data was publicly available and acquired through Fish Pool.

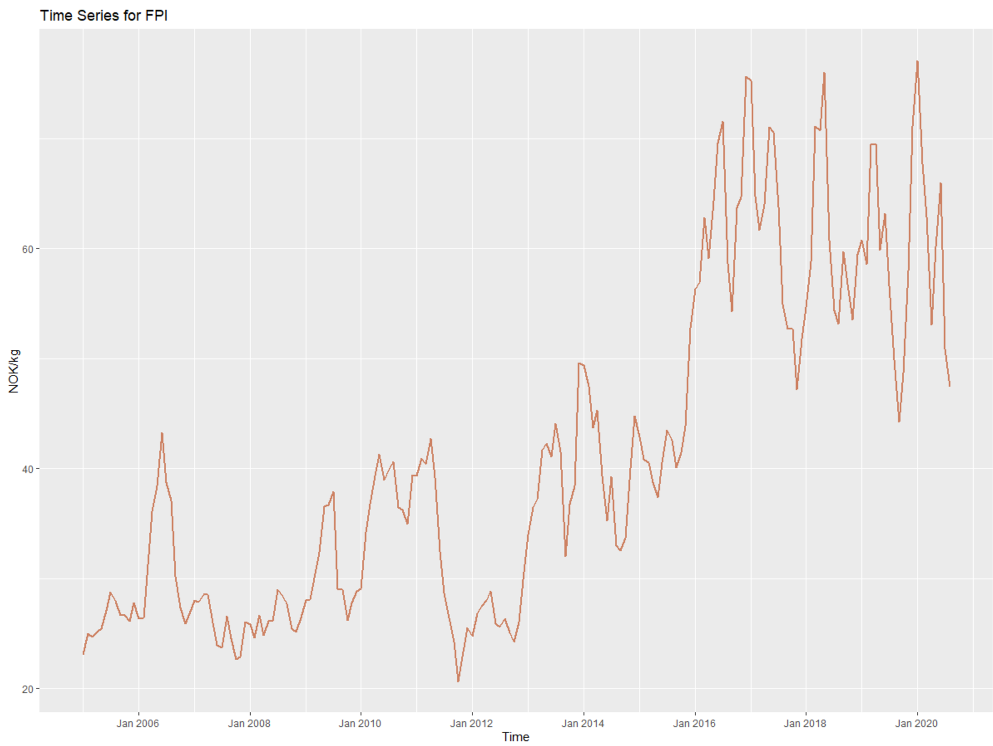


*Figure 5.1: Monthly FPI from January 2005 to August 2020.*

A closer look at Figure 5.1 reveals large fluctuations in price during the period of January 2005 - August 2020. Firstly, there was a large spike in 2006. It is thought that the new regulation in 2005 restricting maximum allowed biomass had a ripple effect, which caused supply shortages into 2006. This happened in conjunction with the introduction of the Fish Pool futures market, making the market more accessible to additional buyers, thereby possibly being responsible for the increase in price. Further on, it is interesting to note during the onset of the financial crisis in late 2007 the price did not change considerably. The reason for this might be attributed to the Chilean sea lice growth, which devastated the production, offsetting any price fall caused by the financial crisis. Equivalently, Norway experienced a similar occurrence of high levels of sea lice between 2010 and 2012 (Lusedata, 2020). This had a detrimental long-term effect on the produced quantity due to premature fish mortality. As a result, the price increased over the period. Additionally, Chile had an algae bloom crisis in 2015-2016, causing the production to fall by 100 000 tonnes (Intrafish, 2019), thereby pushing the price to new heights. Sea lice and algae bloom showcase how companies have to stay vigilant against biological constraints. In the Fall of 2014, the price fell sharply. One reason might be the import ban on salmon Russia introduced in response to sanctions (Salmon Business, 2019). Parts of this ban was circumvented through Belarus. However, this was also shut down in January 2020 due to alleged safety concerns from the Russian Government (Salmon Business, 2019), which is reflected in the sharp fall at the same time. Another event that affected the salmon price fall in 2020 is the outbreak of the SARS-CoV-2 virus, often referred to as Covid-19-/Coronavirus (Holland, 2020). Major shares of the salmon demand come from exporting to foreign retailers and food-service markets, which experienced a significant fall in demand due to national lockdown restrictions. As a consequence, the salmon price has seen a fall. These events highlight how biological constraints, political decisions, and global issues can incur large ramifications on the Norwegian salmon farming industry.

## 5.2 Data and Pre-processing

The data for this thesis is collected through publicly available sources and supplemented with data provided by Kontali Analyse AS. A more detailed data source description can be found in Appendix A.1 Data Sources. The dataset consists of one response variable and 20 features encompassing the timeframe January 2005 to August 2020. This entails 188 observations for each variable. However, when augmenting the data with lagged features, the dataset encompassed 55 predictors. The data was gathered from a plethora of sources in different formats. Therefore, several initial steps were needed to ensure that the data was applicable for usage in both univariate and multivariate forecasting methods. For instance, the frequency of the data was for some variables higher, such as in daily and hourly observations. For the futures prices, we weighted the average based on the volume and price of the contracts made. The exchange rates were given as hourly data, however, this was averaged out throughout the month. The remaining data was provided in a monthly format. There are only two instances of value zero, which was promptly increased to one for transformational and modelling purposes[2].

# 5.3 Descriptive Statistics

## 5.3.1 The FPI

From Figure 5.1 above, it is possible to determine that the FPI sees increased volatility with time. Thus, a transformation was necessary for the FPI. Firstly, a logarithmic transformation was used to reduce the variation. Thereafter, the time series was shown to be non-stationary from KPSS- and ADF-tests. As a result, the FPI had to be first differenced, hereafter denoted as log differenced FPI. When visually inspecting the data after differencing, it looked to be stationary. This is shown in Figure 5.3, which appears to exhibit stationarity. However, the autocorrelation function in Figure 5.2 seems to show a positive correlation at every 12 lags and a negative correlation at every 6 + 12 lags. In essence, a visual inspection was not adequate to confirm the stationarity of the FPI.

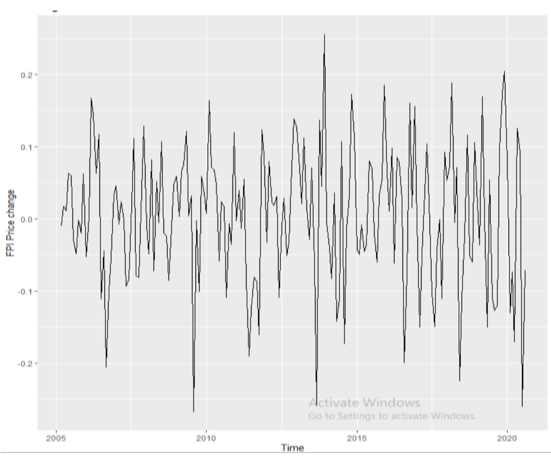


*Figure 5.3: Log differenced FPI.*

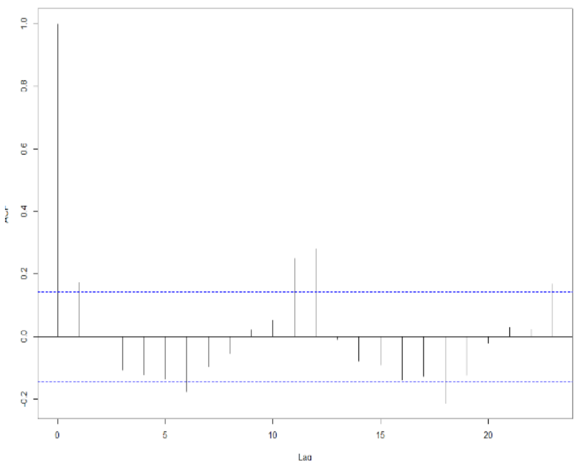


*Figure 5.2: Autocorrelation plot of FPI after Log transformation and differencing.*

[2] Smolt release was modified January -2006 and -2011. The change had negligible effect on the models.

To verify the stationarity in an objective way, the time series was validated through unit root tests. Firstly, KPSS tests showed the test statistic with four lags to be lower than the critical value of 0.216, corresponding to significance at the 1% level. Thus, $h_0$ is not rejected, and stationarity is validated at the 1% significance. Following an ADF-test showed the test statistic for five lags lower than the critical value, -7.278 < - 3.17. The H0 was rejected, and the test validated at 1% significance, further underscoring the stationarity of the log differenced FPI. In essence, two tests were used as a safeguard to ensure stationary, and the root tests imply strong evidence of stationarity in the FPI log differenced time series.

### 5.3.2 The Features

All the explanatory variables show non-stationarity from the original data, in addition to the smolt release and sea temperature, which show annual cycles (i.e., seasonal pattern). This was validated both visually (see Figure 5.4) and through the root tests.

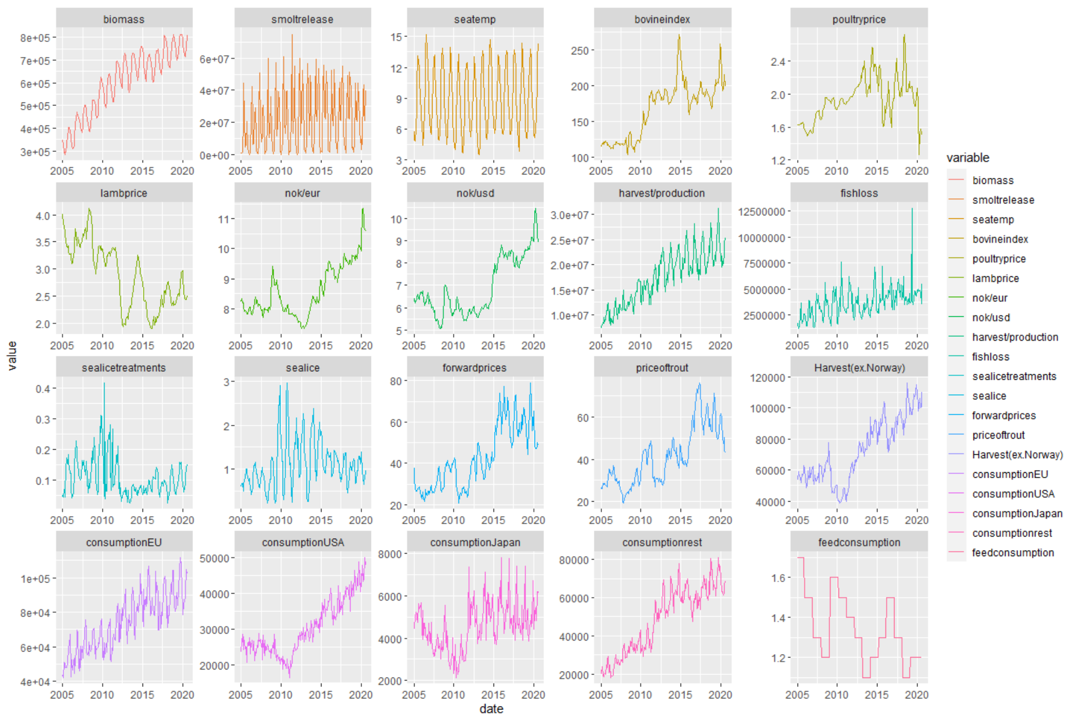


*Figure 5.4: Graphs of all explanatory variables from January 2005 to August 2020.*

Therefore, log transformation was performed along with first differencing of all variables. Consequently, the unit root tests were performed on the log differenced data to ensure stationarity. Firstly, the KPSS test showed that every variable displayed stationarity at 1%

significance following the log differencing. The ADF test resulted in similar outcomes, with every feature showing stationarity at the 1% level, with the exception of lamb prices, which appear stationary at 5%. This might be the result of highly volatile lamb prices over the period. The tests appear to show enough evidence of stationarity before we begin forecasting.

*Table 5.1 Descriptive statistics and root test values for all log differenced data.*

| Feature Information | Descriptive Statistics | | | | Root Tests | |
|---|---|---|---|---|---|---|
| *Variables* | *Observations* | *Std. deviation* | *Mean* | *FPI Correlation* | *KPSS* | *ADF* |
| Fish Pool Index (NOK/kg) | 186 | 0.244 | 0.054 | 1.000 | 0.086 | -7.278 |
| Biomass (Tonnes, Norway) | 186 | 0.063 | 0.056 | -0.337 | 0.110 | -9.618 |
| Smolt release (#Individual Salmon, Norway) | 186 | 0.909 | 0.117 | 0.019 | 0.072 | -10.843 |
| Sea Temperature (Celcius, Norway) | 186 | 0.122 | 0.001 | -0.138 | 0.048 | -12.695 |
| Bovine Index (Index) | 186 | 0.141 | 0.040 | -0.267 | 0.130 | -6.763 |
| Poultry Price (USD/kg) | 186 | 0.121 | 0.003 | -0.211 | 0.104 | -5.593 |
| Lamb Pice (USD/kg) | 186 | 0.198 | -0.019 | -0.080 | 0.087 | -4.759 |
| NOK/EUR | 186 | 0.062 | 0.019 | 0.128 | 0.111 | -5.111 |
| NOK/USD | 186 | 0.116 | 0.026 | 0.159 | 0.093 | -4.819 |
| Feed Consumption | 186 | 0.172 | 0.044 | -0.421 | 0.153 | -7.478 |
| Harvest (#Individuals, Norway) | 186 | 0.109 | 0.062 | -0.428 | 0.049 | -6.994 |
| Fish Loss (#Individual Salmon, Norway) | 186 | 0.299 | 0.055 | -0.072 | 0.062 | -7.916 |
| Sea Lice Treatments (Share of Salmon) | 186 | 0.450 | 0.003 | 0.107 | 0.160 | -8.672 |
| Sea Lice (Average amount/Salmon) | 186 | 0.412 | 0.020 | -0.132 | 0.053 | -8.700 |
| Forward Prices (NOK/kg) | 186 | 0.230 | 0.049 | -0.163 | 0.109 | -6.812 |
| Price of Trout (NOK/kg) | 186 | 0.248 | 0.041 | 0.816 | 0.099 | -6.689 |
| Harvest (Tonnes, ex. Norway) | 186 | 0.163 | 0.043 | -0.406 | 0.146 | -7.242 |
| Consumption EU (Tonnes) | 186 | 0.080 | 0.041 | -0.576 | 0.088 | -7.166 |
| Consumption USA (Tonnes) | 186 | 0.120 | 0.039 | -0.454 | 0.162 | -7.120 |
| Consumption Japan (Tonnes) | 186 | 0.236 | 0.003 | -0.232 | 0.162 | -6.843 |
| Consumption Other (Tonnes) | 186 | 0.125 | 0.077 | -0.615 | 0.111 | -7.581 |

*ADF*: $\tau_\alpha = 5\% < -3.45$; $\tau_\alpha = 1\% < -4.04$

*KPSS*: $\tau_\alpha = 10\% > 0.119$; $\tau_\alpha = 5\% > 0.146$; $\tau_\alpha = 1\% > 0.216$

# 5.4 Lag Feature Analysis

In chapter 3, a thorough analysis of the important features and their lag structure was presented. However, a closer analysis is required to confidently understand the lag structure between the log differenced univariate and log differenced explanatory variables. This is achieved through a cross-correlation function (see section 4.3).

The cross-correlation function permits for an improved overview of the lag relationship between the FPI and explanatory variables. However, we concluded that results that show significant discrepancies from what was advised by MOWI (2020), industry experts, and research from chapter 3 would be discarded. It is also important to notice that many of the explanatory variables display both positive and negative cross-correlation with the FPI. This is perhaps a consequence of the seasonality of the variables. Thus, our analysis would resolve this by using the assumptions presented in chapter 3 with regard to the correlational impact of each predictor. However, this is only to better understand the features seen as the directional relationship has no impact on the modelling, only the lags. There are several variables that showed lag structures not in line with our assumptions. For instance, lamb price displayed a high positive correlation at 10 months, which is not consistent with the reasoning from chapter 3 or the other cross-correlation returns from alternative proteins. Hence, it is not in concurrence with industry assumptions. The same applies to the exchange rate NOK/USD, which showed positive and negative correlation at lag 11, a far difference from lag six and 12 advised earlier. Additionally, sea lice treatment showed both positive correlations over the lags two to four and 11-14, which also somewhat contradicts the assumption from chapter 3, which were that sea lice has a negatively short-term effect and a positive long-term effect on the FPI. The rest were mostly in line with industry assumptions, however, some showed minor alterations in the lag structure. In essence, chapter 3 lays the foundation for the proposed lag structure. However, the cross correlational analysis provides valuable insight for industry participants and as a tool to modify lags structures as long as it is reinforced by an underlying relationship presented in chapter 3.

*Table 5.2: Lag structure.*

| **Feature Information** | **Lag Structure** | | |
|---|---|---|---|
| | Chapter 3 | CCF Proposed | Final |
| Biomass (Tonnes, Norway) | 4, 8 | 5 - 8 | 5 - 8 |
| Smolt release (#Individual Salmon, Norway) | 18 - 24 | 17, 18 | 18 - 24 |
| Sea Temperature (Celcius, Norway) | 3 - 5 | 4 - 7 | 4 - 7 |
| Bovine Index (Index) | 1 | 1 - 2 | 1-2 |
| Poultry Price (USD/kg) | 1 | 1 | 1 |
| Lamb Pice (USD/kg) | 1 | 10 | 1 |
| NOK/EUR | 6, 12 | 1 | 6, 12 |
| NOK/USD | 6, 12 | 11 | 6, 12 |
| Feed Consumption | 2 - 4 | 4 - 8 | 4 - 6 |
| Harvest (#Individuals, Norway) | 1 | 1 | 1 |
| Fish Loss (#Individual Salmon, Norway) | 1 | 4 | 1 |
| Sea Lice Treatments (Share of Salmon) | 3, 12 | 5, 9 | 3, 12 |
| Sea Lice (Average amount/Salmon) | 3, 12 | 5, 6 | 3, 12 |
| Forward Prices (NOK/kg) | 1 - 4 | 4 | 1 - 4 |
| Price of Trout (NOK/kg) | 1 | 1 | 1 |
| Harvest (Tonnes, ex. Norway) | 1 | 1 | 1 |
| Consumption EU (Tonnes) | 1 | 1 | 1 |
| Consumption USA (Tonnes) | 1 | 1 | 1 |
| Consumption Japan (Tonnes) | 1 | 1 | 1 |
| Consumption Other (Tonnes) | 1 | 1 | 1 |

# 6 Forecasting and Analysis

In this section, we will look into the forecasting analysis and results from the models elaborated in chapter 4. Further, we will evaluate the accuracy of each model and compare them against each other to find the best method.

## 6.1 Univariate Methods

For univariate forecasting models, we have chosen the ETS and the ARIMA model, as mentioned earlier. These models base their forecast on historical observations of the forecast variable itself. Therefore, we start with creating a time series containing all observations of this variable using the ts function in R. This function uses the argument frequency = 12 to specify it as monthly data, and the start year and month with the start argument. Further, we divide the time series into two separate time series; training set and testing set, where the training set is used to train the forecasting model with rolling-origin (see section 4.4.3), while the test set is used to evaluate the model accuracy.

### 6.1.1 Benchmarks

Before we start to build our ETS and ARIMA models, we want to create some simple forecasting models as benchmarks (see section 4.1.1). To evaluate these forecasts against each other, we use the accuracy function in R, which calculates a number of accuracy measurements. As mentioned earlier, we will use RMSE and MAE as our accuracy measurements. The accuracy is measured for 1 to 12 months ahead. The best forecast accuracy for each period will become the forecast benchmark.

*Table 6.1: RMSE and MAE accuracy measure for naïve and rwdrift.*

| | naïve RMSE | naïve MAE | rwdrift RMSE | rwdrift MAE |
|---|---|---|---|---|
| 1 month | 2.417 | 2.417 | 2.602 | 2.602 |
| 2 months | 4.529 | 4.174 | 4.820 | 4.451 |
| 3 months | 3.947 | 3.580 | 4.287 | 3.949 |
| 4 months | 4.754 | 4.336 | 4.732 | 4.429 |
| 5 months | 7.396 | 6.176 | 7.051 | 6.065 |
| 6 months | 10.209 | 8.272 | 9.661 | 7.995 |
| 7 months | 10.254 | 8.593 | 9.600 | 8.171 |
| 8 months | 9.641 | 7.864 | 8.992 | 7.310 |
| 9 months | 9.406 | 7.796 | 8.983 | 7.489 |
| 10 months | 8.923 | 7.034 | 8.546 | 6.942 |
| 11 months | 8.812 | 7.086 | 8.320 | 6.818 |
| 12 months | 8.598 | 6.975 | 8.291 | 6.913 |

From Table 6.1, we can observe all accuracy measures for both benchmark models. The average RMSE and MAE for all forecasts of the naïve method is 7.407 and 6.192, and the rwdrift method has an average RMSE and MAE of 7.157 and 6.094. In terms of directional accuracy, the naïve and rwdrift methods predict with 75% accuracy. We will now continue to build an ETS and ARIMA forecasting model and then evaluate those model accuracies against these benchmarks.

### 6.1.2 ETS Forecast

In R, we use the ets function to find a suitable ETS model, and by using the argument model = "ZZZ", the function automatically selects the best model for the given time series. This R function suggests that ETS (M, Ad, M) is the most suitable model, which is an ETS model with a multiplicative error component, an additive damped trend component, and a multiplicative seasonal component.

Now, as we have selected an appropriate ETS model, we want to forecast with this model for different periods. To create the forecast itself, we use the R function forecast.ets, then select the ETS model we have chosen and the number of months ahead we want to forecast. Just as the benchmarks, we forecast the ETS model for 1 to 12 months ahead. To evaluate the forecasts, we start with a visual inspection by using the autoplot function and supply with the autolayer function to place a layer of the training set and the test set. From Figure 6.1, we are able to see how the ETS (M,Ad,M) forecast model performs compared to the test set. We see that the forecast model catches the seasonal pattern but could not reproduce the test set identically. However, all observations from the test set are within the 80 % confidence interval.

To measure the accuracy, we stored the forecasting result in a matrix and compared it against the test set we created in the beginning with the accuracy function in R.

*Table 6.2: RMSE and MAE accuracy measures from ETS.*

| | ETS RMSE | ETS MAE |
|---|---|---|
| 1 month | 3.764 | 3.764 |
| 2 months | 5.012 | 4.885 |
| 3 months | 4.093 | 3.280 |
| 4 months | 5.688 | 4.684 |
| 5 months | 8.967 | 7.050 |
| 6 months | 12.239 | 9.589 |
| 7 months | 12.299 | 10.027 |
| 8 months | 11.656 | 9.437 |
| 9 months | 11.146 | 9.008 |
| 10 months | 10.581 | 8.227 |
| 11 months | 10.421 | 8.266 |
| 12 months | 10.121 | 8.069 |

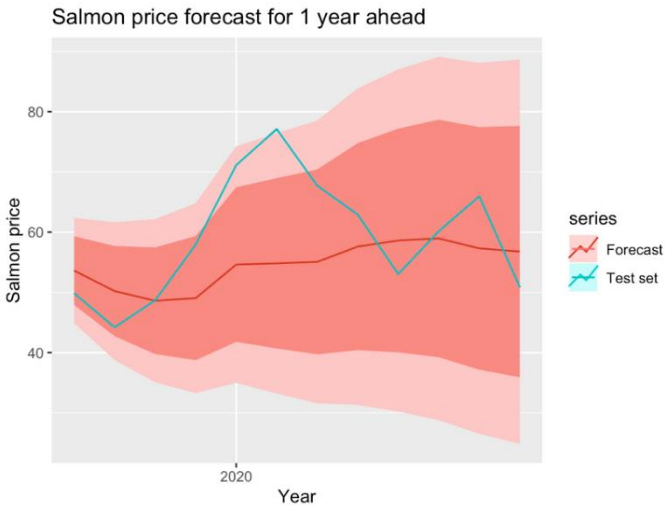

*Figure 6.1: ETS forecast with 80% and 95% confidence interval compared to the test set.*

From Table 6.2, we can see that the ETS model performs very well, especially on a short-term basis. At the lowest, the RMSE and MAE is measured to be 3.763 for a 1 month ahead forecast and an RMSE and MAE of 12.299 and 10.027 at the highest at 7 months ahead forecast. The average RMSE and MAE can be calculated to 8.832 and 7.190. Directional prediction accuracy is calculated to be just below 60%. Further, we will build an ARIMA forecasting model to see if it can outperform the ETS model and the other benchmarks.

### 6.1.3 ARIMA Forecast

As mentioned earlier, an ARIMA forecast requires the data to be stationary. In section 5.3.1, we concluded that we had to apply log transformation and first differencing in order to obtain a stationary time series. In R, we use the log and diff functions to apply log transformation and first differencing.

Now that the time series is stationary, we start with the automatic alternative and use this model to forecast and calculate the forecast accuracy, then use the manually chosen model and use this model to forecast, and lastly, compare these models against each other. For the automatic solution, we use the auto.arima function in R, and as arguments we set ic =AICc, which sets AICc as the information criteria, then d = 1 and D = 0, which sets the number of first difference equal one and seasonal difference equal zero. This function returns the ARIMA model ARIMA (0,1,0) (0,0,1). Now, as the model is defined, we continue to forecast with this model by using the forecast function in R and sets h=12 for a forecast over 12 months. To evaluate this forecast, we start by visually inspecting the plot by using the autoplot function in R and autolayer functions to add layers of the test set and train set. Further, we use RMSE and MAE as our accuracy measurements by using the accuracy function in R.

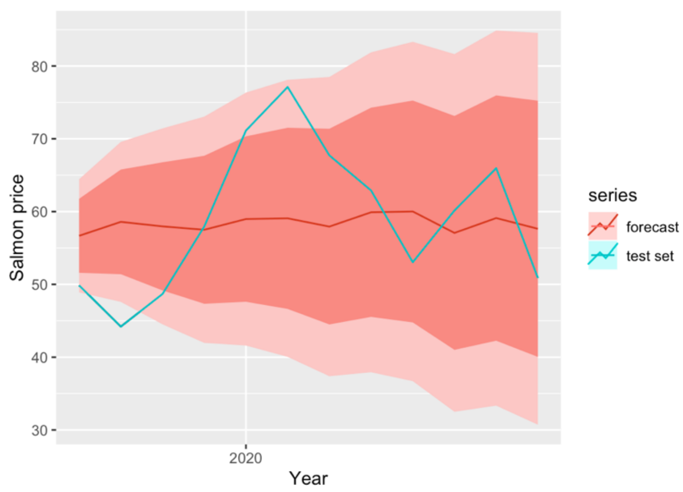


*Figure 6.2: Auto ARIMA forecast with 80 % and 95% confidence interval.*

*Table 6.3: RMSE and MAE for automatically generated ARIMA model.*

| | ARIMA RMSE | ARIMA MAE |
|---|---|---|
| 1 month | 6.813 | 6.813 |
| 2 months | 11.255 | 10.599 |
| 3 months | 10.640 | 10.162 |
| 4 months | 9.217 | 7.732 |
| 5 months | 9.876 | 8.617 |
| 6 months | 11.640 | 10.187 |
| 7 months | 11.394 | 10.129 |
| 8 months | 10.710 | 9.238 |
| 9 months | 10.362 | 8.986 |
| 10 months | 9.878 | 8.396 |
| 11 months | 9.643 | 8.255 |
| 12 months | 9.436 | 8.131 |

As Table 6.3 displays, the ARIMA model generated by the auto.arima function has a quite high RMSE and MAE. At the lowest, the ARIMA (0,1,1) (0,0,1) has an RMSE and MAE of 6.813 for one month ahead forecast. However, for medium time periods such as 6 months ahead forecast, we see that the RMSE and MAE are as high as 11.640 and 10.186. The average RMSE and MAE are calculated to be 10.072 and 8.937. In addition, the directional accuracy is calculated to be 50%. Further, we want to look at how the manually built ARIMA model performs and compare it against the forecast accuracy of the automatically created model.

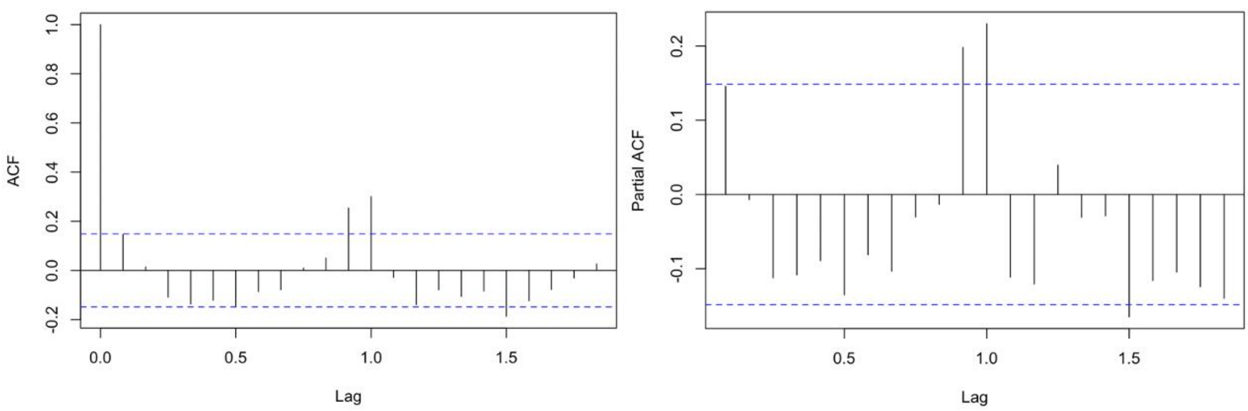


*Figure 6.3: Autocorrelation plot (left) and Partial autocorrelation plot (right).*

As mentioned in section 4.1.3, we use the ACF plot and PACF plot to determine the parameter p, P, q, and Q in the manually created ARIMA model. In Figure 6.3, we see that there are three significant positive lags and one significant negative lags in the ACF plot, and two significant positive and one significant negative lags in the PACF plot. This implies that our ARIMA model should be ARIMA (2,1,3) (1,0,1). The information criteria AICc is used to evaluate how good the values are for p, P, q, and Q. By using the Arima function in R, setting the order argument to be (2,1,3) and the seasonal argument to be (1,0,1), we can get a summary of how well this model performs. This model returns an AICc of -357.14 and will be our ARIMA model benchmark. Further, we want to see if we can find a better model with lower AICc. We will try to increase/decrease the values of p, P, q, and Q. The following models were tested by a similar approach as our benchmark ARIMA model:

*Table 6.4: Overview of AICc tested ARIMA models.*

| Model number | ARIMA Model | AICc |
|---|---|---|
| Model 0 | ARIMA (3,1,3) (1,0,1) | - 365.38 |
| Model 1 | ARIMA (4,1,3) (1,0,1) | - 363.26 |
| Model 2 | ARIMA (3,1,2) (1,0,1) | - 363.31 |
| Model 3 | ARIMA (3,1,4) (1,0,1) | - 360.50 |
| Model 4 | ARIMA (3,1,3) (0,0,1) | - 362.97 |
| Model 5 | ARIMA (3,1,3) (2,0,1) | - 363.97 |
| **Model 6** | **ARIMA (3,1,3) (1,0,0)** | **- 366.17** |
| Model 7 | ARIMA (3,1,3) (1,0,2) | - 364.15 |

As displayed in Table 6.4, model 6, ARIMA (3,1,3) (1,0,0), has an AICc of -366.17, which is even lower than our ARIMA benchmark model and would therefore replace the ARIMA benchmark model.

Now that we have built an appropriate ARIMA model, we want to forecast with this model. Similar to the benchmark- and ETS models, we forecast for 1 to 12 months. In R, we use the forecast function and set the argument h equal to the period we want to forecast. To evaluate the forecast, we start with a visual inspection of the forecast by using the autoplot function and add layers of the train set and test the set with the autolayer function in R.

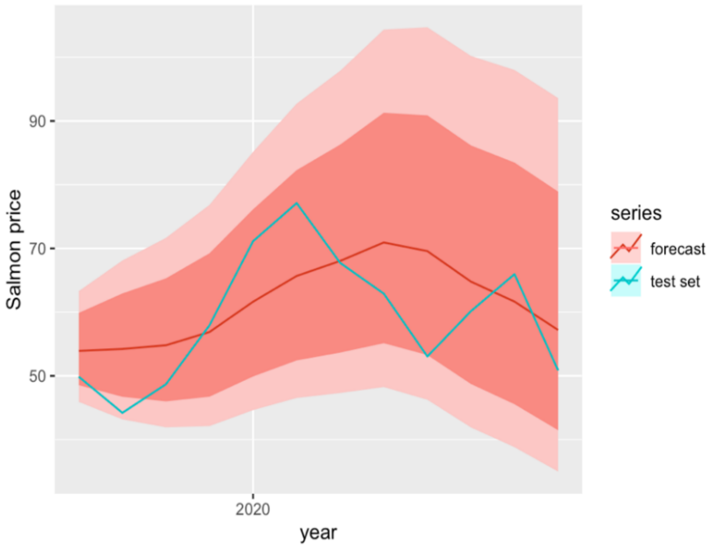

Figure 6.4: Graph of the ARIMA forecast.

Table 6.5: MAE and RMSE for optimal ARIMA model.

| | ARIMA RMSE | ARIMA MAE |
|---|---|---|
| 1 month | 4.051 | 4.051 |
| 2 months | 7.641 | 7.035 |
| 3 months | 7.170 | 6.730 |
| 4 months | 6.232 | 5.312 |
| 5 months | 7.008 | 6.149 |
| 6 months | 7.919 | 7.030 |
| 7 months | 7.333 | 6.071 |
| 8 months | 7.424 | 6.317 |
| 9 months | 8.909 | 7.452 |
| 10 months | 8.578 | 7.170 |
| 11 months | 8.281 | 6.910 |
| 12 months | 8.137 | 6.862 |

We can from Figure 6.4 see that the ARIMA model forecast is adequate, and all observations from the test set are within the 95% confidence interval. To evaluate the forecast accuracy, we use the R function accuracy and compare it against the test set, as we did with the other models. From  with an RMSE accuracy measure for all periods, we can see that the manually selected ARIMA model has an overall lower forecast error. The lowest RMSE and MAE are calculated for 1 month ahead forecast at 4.050.  This ARIMA model has an average RMSE and MAE of 7.390 and 6.424, which are both lower than the calculated average forecast error for the automatically picked model. The directional accuracy is calculated to be just below 60%. However, to determine if any of the models are better in different time period intervals, we divide the time periods short, medium, and long term. Here, the short term represents 1 to 3 months, the medium-term represents 4 to 7 months, and the long term represents 8 to 12 months.

Table 6.6: RMSE and MAE for different horizons.

| | **Automatically** | | **Manually** | |
|---|---|---|---|---|
| | RMSE | MAE | RMSE | MAE |
| Short | 9.569 | 9.192 | 6.288 | 5.939 |
| Medium | 10.532 | 9.166 | 7.123 | 6.141 |
| Long | 10.006 | 8.601 | 8.266 | 6.942 |

From Table 6.6, we can see that the automatically created ARIMA model has a higher RMSE and MAE than the manually created ARIMA model for all time period intervals. Based on this, we chose to discard the automatically created ARIMA model and continue our analysis with the manually selected model, ARIMA (3,1,3) (1,0,0). Further, we want to compare this ARIMA model with the ETS model we created earlier and evaluate which model performs better.

### 6.1.4 Univariate Method Comparison

As mentioned earlier, the automatically selected ARIMA model has an average RMSE and MAE of 7.390 and 6.424. In other words, the ARIMA model has a lower average forecast error than the ETS model, with an average RMSE and MAE of 8.832 and 7.190. However, if we look at the lowest RMSE and MAE values of the ETS models, it is lower than the forecast error measured for the ARIMA model for all periods. This shows that the ETS model has a high variation in forecast accuracy and may perform better than the ARIMA models in certain time periods. To further investigate how the ARIMA and ETS model performs in the different time periods, we merge the two tables into one and split it up by RMSE and MAE. Additionally, we have added conditional formatting for visual simplicity. The colour green indicates the best model, and red indicates the worst model. We also created a matrix divided into short-, medium- and long-term forecasting as earlier to evaluate the accuracy for given time intervals.

*Table 6.7: ETS and ARIMA comparison 12 steps ahead.*

| | RMSE | | MAE | |
|---|---|---|---|---|
| | ETS | ARIMA | ETS | ARIMA |
| 1 month | 3.764 | 4.051 | 3.764 | 4.051 |
| 2 months | 5.012 | 7.641 | 4.885 | 7.035 |
| 3 months | 4.093 | 7.170 | 3.280 | 6.730 |
| 4 months | 5.688 | 6.232 | 4.684 | 5.312 |
| 5 months | 8.967 | 7.008 | 7.050 | 6.149 |
| 6 months | 12.239 | 7.919 | 9.589 | 7.030 |
| 7 months | 12.299 | 7.333 | 10.027 | 6.071 |
| 8 months | 11.656 | 7.424 | 9.437 | 6.317 |
| 9 months | 11.146 | 8.909 | 9.008 | 7.452 |
| 10 months | 10.581 | 8.578 | 8.227 | 7.170 |
| 11 months | 10.421 | 8.281 | 8.266 | 6.910 |
| 12 months | 10.121 | 8.137 | 8.069 | 6.862 |

Table 6.8: ETS and ARIMA comparison for short, medium, and long horizons.

| | ARIMA | | ETS | |
|---|---|---|---|---|
| | RMSE | MAE | RMSE | MAE |
| Short | 6.288 | 5.939 | 4.290 | 3.976 |
| Medium | 7.123 | 6.141 | 9.798 | 7.838 |
| Long | 8.266 | 6.942 | 10.785 | 8.601 |

From Table 6.7, we can observe that there is a clear pattern in accuracy. The ARIMA model performs better than the ETS overall with a lower RMSE and MAE in 8 out of 12 time periods, which is also reflected by lower average RMSE and MAE of the ARIMA model. And it is clear that the ETS model is better in the short term, but after 4 months ahead forecast and onwards, the ARIMA model has the lowest forecast error. This assumption can be confirmed in table 6.8. By looking into time period intervals, it is possible to observe that the ETS models perform better than the ARIMA model in the short term but worse in a medium to long term forecast and therefore has a higher RMSE and MAE for those time period intervals. The question is now if any of these models perform better than the benchmark models we created in the beginning. To compare the models against each other, we add all four forecasting models into one table. For visual simplicity, we add some formatting to the matrix. The formatting is colour based on a colour scale from green (best) to red (worst).

Table 6.9: Univariate model comparison.

| | RMSE | | | | MAE | | | |
|---|---|---|---|---|---|---|---|---|
| | ETS | ARIMA (M) | Naïve | Rwdrift | ETS | ARIMA (M) | Naïve | Rwdrift |
| 1 month | 3.764 | 4.051 | 2.417 | 2.602 | 3.764 | 4.051 | 2.417 | 2.602 |
| 2 months | 5.012 | 7.641 | 4.529 | 4.820 | 4.885 | 7.035 | 4.174 | 4.451 |
| 3 months | 4.093 | 7.170 | 3.947 | 4.287 | 3.280 | 6.730 | 3.580 | 3.949 |
| 4 months | 5.688 | 6.232 | 4.754 | 4.732 | 4.684 | 5.312 | 4.336 | 4.429 |
| 5 months | 8.967 | 7.008 | 7.396 | 7.051 | 7.050 | 6.149 | 6.176 | 6.065 |
| 6 months | 12.239 | 7.919 | 10.209 | 9.661 | 9.589 | 7.030 | 8.272 | 7.995 |
| 7 months | 12.299 | 7.333 | 10.254 | 9.600 | 10.027 | 6.071 | 8.593 | 8.171 |
| 8 months | 11.656 | 7.424 | 9.641 | 8.992 | 9.437 | 6.317 | 7.864 | 7.310 |
| 9 months | 11.146 | 8.909 | 9.406 | 8.983 | 9.008 | 7.452 | 7.796 | 7.489 |
| 10 months | 10.581 | 8.578 | 8.923 | 8.546 | 8.227 | 7.170 | 7.034 | 6.942 |
| 11 months | 10.421 | 8.281 | 8.812 | 8.320 | 8.266 | 6.910 | 7.086 | 6.818 |
| 12 months | 10.121 | 8.137 | 8.598 | 8.291 | 8.069 | 6.862 | 6.975 | 6.913 |

The first thing we should notice from Table 6.9 is that none of the benchmark models, naïve and rwdrift, are the least accurate model in one single time period. Both naïve and rwdrift has a low forecast error for lower time periods but still has decent forecast accuracy for longer time periods. The ETS model does not perform as the best forecast model in one single time period and is therefore discarded. When it comes to the ARIMA model, it is a little more difficult to judge. The benchmark models perform well in the short term, but after 4-5 months ahead forecast, we see that the ARIMA model has the lowest forecast error in almost every forecast period. Therefore, it is reasonable to conclude that for shorter time periods, the benchmark models as the naïve STL decomposed model performs the best, but for longer time periods, the ARIMA (3,1,3) (1,0,0), which is the manually created model, is the best forecasting model. Further, we want to investigate if we can build a better model by introducing other explanatory variables.

## 6.2 Multivariate Methods

The structure of the multivariate modelling is composed in the same order, which was presented in chapter 4. Firstly, all relevant components to modelling 12-step ahead forecast are presented, the results, and the variable subset selection chosen for the respective models. After accounting for lags and first differencing, more than 25 observations are omitted, while 12 observations are left out in the test set. Consequently, the modelling is based on 151 observations ranging from February 2007 to August 2019. Simply put, the computing entailed creating a user-defined wrapper function for modelling followed by a user-defined prediction function.

### 6.2.1 LASSO

The following section is used to illustrate and discuss the modelling across different forecast horizons and to combine these into an optimal 12 step ahead forecast for the FPI using LASSO. Furthermore, a closer deliberation into the variable selection is conducted to better understand the features which perform optimally when using LASSO.

Firstly, we trained the data for each separate horizon using nested cross-validation (see section 4.4.3) in order to find the lambda value, which minimizes the loss function and thereby creates the best fit.

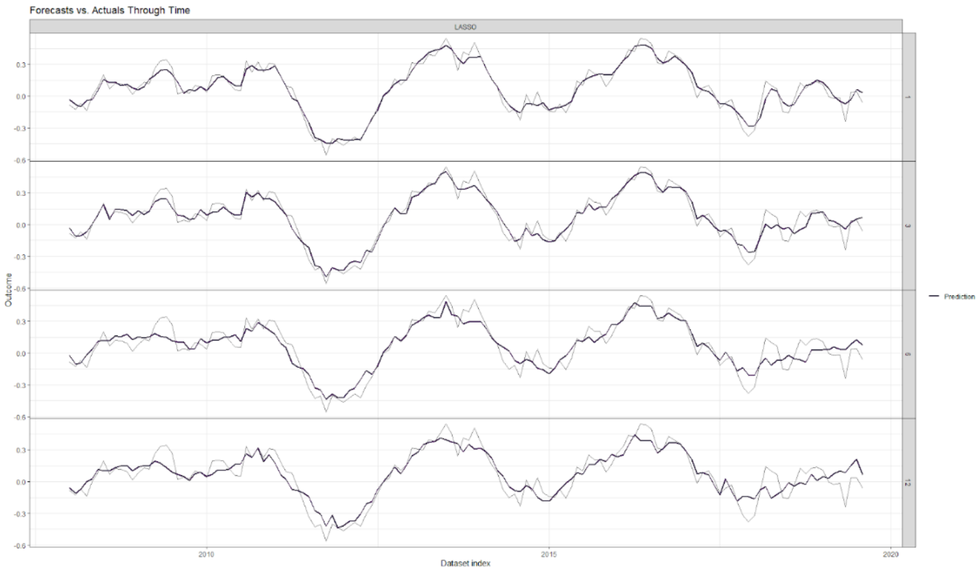


*Figure 6.5: The fit for horizons 1-, 3-, 6-, and 12-steps ahead of the log difference.*

Visually it is difficult to determine which horizon results in a superior fit, however using the metrics deliberated on in section 4.4.2, it is simpler to illustrate that according to both MAE and RMSE that horizon 1 step ahead shows the closest fit, followed by 6-, 3-, and 12-steps ahead.

*Table 6.10: MAE and RMSE of the log differenced data for horizons 1, 3, 6, and 12.*

| Forecast Horizon | Window Start | Window Stop | MAE | RMSE |
|---|---|---|---|---|
| 1 | 01.02.2007 | 01.08.2019 | 0.042 | 0.054 |
| 3 | 01.02.2007 | 01.08.2019 | 0.066 | 0.084 |
| 6 | 01.02.2007 | 01.08.2019 | 0.061 | 0.077 |
| 12 | 01.02.2007 | 01.08.2019 | 0.067 | 0.084 |

Naturally, computing different h steps ahead forecasts showed different results for the metrics compared to the fit.

*Table 6.11: Errors of log differenced LASSO forecast 1-, 3-, 6-, and 12 steps ahead.*

| Forecast Horizon | MAE | RMSE |
|---|---|---|
| 1 | 0.103 | 0.103 |
| 3 | 0.120 | 0.122 |
| 6 | 0.115 | 0.122 |
| 12 | 0.127 | 0.136 |

For instance, the error rate from the out-of-sample measured in MAE and RMSE was significantly higher for all h steps ahead. The order of error rate among the horizons was equivalent in MAE. RMSE, however, measured increasing error with the number of forecasting steps ahead. This was anticipated given the unpredictable nature of forecasting, and it inherently becoming harder to determine forecast with the advancement of time. As a result, the further steps ahead forecast will almost always become increasingly error-prone.

The direct 12-steps ahead forecast showed unfavourable results. Therefore, a different approach was used to optimize the 12-step ahead forecast. This was done by combining the different horizons in order to create an optimized 12 step ahead forecast based on the fit and forecast of the different horizons. In addition, all combined results were reverted into the original data structure (i.e., regress the difference and log steps) in order to increase interpretability. Therefore, the results will be comparable to the FPI and not the log differenced FPI. Hereafter only reverted data will be used to illustrate forecast out-of-sample.

$$F_{C12} = F_{1,1} + F_{3,2-3} + F_{6,4-6} + F_{12,7-12} \tag{6.1}$$

From equation 6.1, $F_{C12}$ represents the combined forecast. Forecast one, $F_{1,1}$, consist of the forecast for the first month, while $F_{3,2-3}$ is the 3-steps ahead forecast, and includes the forecast for 2 and 3 months ahead, $F_{6,4-6}$ and $F_{12,7-12}$ follows a similar pattern.

*Table 6.12: MAE and RMSE forecast error for the combined 12-step ahead prediction.*

| Forecast Horizon | MAE | RMSE |
|---|---|---|
| 1 | 2.339 | 2.339 |
| 2 | 2.452 | 2.454 |
| 3 | 5.122 | 6.364 |
| 4 | 8.901 | 11.522 |
| 5 | 11.812 | 14.705 |
| 6 | 11.958 | 14.389 |
| 7 | 11.207 | 13.560 |
| 8 | 10.164 | 12.725 |
| 9 | 9.615 | 12.123 |
| 10 | 9.828 | 12.085 |
| 11 | 9.158 | 11.547 |
| 12 | 8.846 | 11.165 |

Table 6.12 shows that the results were still inadequate, missing the actual forecast significantly. This is apparent, knowing the price varied between NOK44-NOK77 in the test set. However, it can be seen that combining 1-, 3-, 6-, and 12-steps ahead into a single one-year forecast creates slightly improved results measured both in MAE and RMSE compared to a direct 12-steps ahead forecast on its own. This can be seen in the Table 6.13 below.

*Table 6.13: MAE and RMSE for combined and direct 12 steps ahead forecast.*

| Forecast Horizon | MAE | RMSE |
|---|---|---|
| Combined 12-step | 8.846 | 11.165 |
| Direct 12-step | 9.367 | 12.316 |

The metrics exhibit that combining forecasts can improve accuracy. For instance, the MAE decreased by 0.5214, which amounts to a 5.57% reduction in error rate. For RMSE, the same results were evident, however, more pronounced at a 9.34% decrease. Furthermore, when inspecting the breakdown for each period for the combined forecast, it can be seen that certain months are quite accurate while in other instances, there is a substantial discrepancy between actual and predicted. The magnitude of the errors is somewhat in line with previous assumptions that earlier periods saw lower error. For instance, the first two months had errors around 2 for both MAE and RMSE. This can be seen to increase for periods 3-8, before slightly falling.

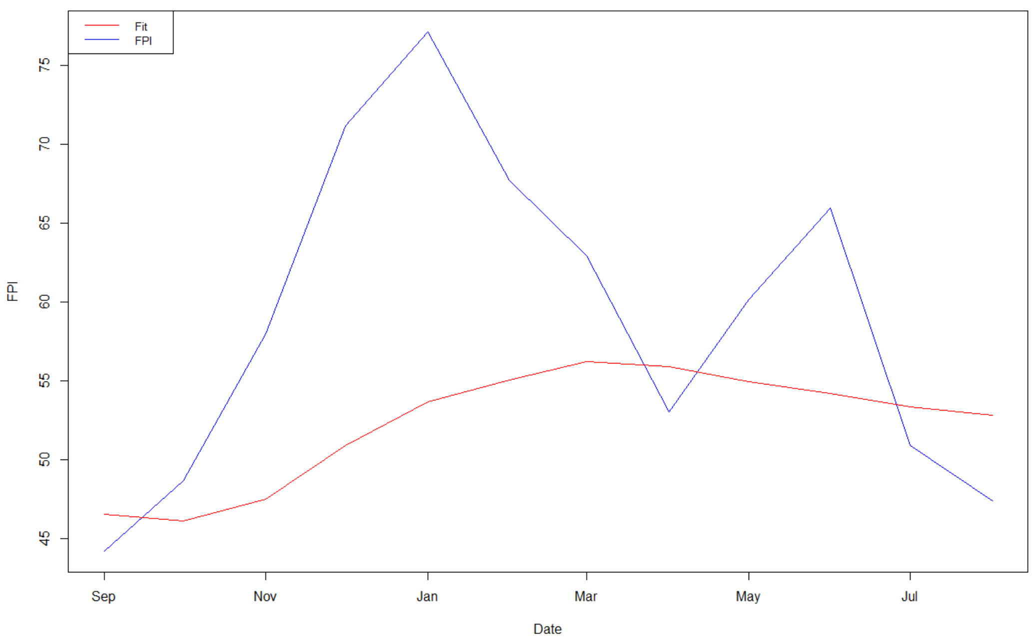


*Figure 6.6: Forecast for combined 12 steps ahead forecast.*

Additionally, the model captures the direction of the price movements inadequately. For the combined 12-steps ahead forecast, the model captures directional movements only 50% of the time, which leaves much room for improvement. Furthermore, the LASSO performed a variable selection, which resulted in eight variables from the 12-steps ahead direct forecast. Among them were smolt release at lags 20 and 24, all protein replacements, sea lice at lag zero and three, sea temperature at lag three, and consumption in Japan. The largest coefficient was the price of trout.

### 6.2.2 GAM

This section deliberates on and discusses the modelling for GAM 12 step-ahead forecasting of the FPI based on the methodology outlined in section 4.2.2. The first step after the cleaning and adding the augmented lagged dataset was the variable selection. This was done through the partly exhaustive forward stepwise selection explained in section 4.2.2.1.

For this approach, the best subset for q number of explanatory variables where chosen. For computational purposes, only subsets of size q ranging 1:20 were allowed (i.e., larger subsets would be infeasible to run). The performance of each subset of q models was then compared using AICc. In the case of AICc, a lower score showed superior model performance.

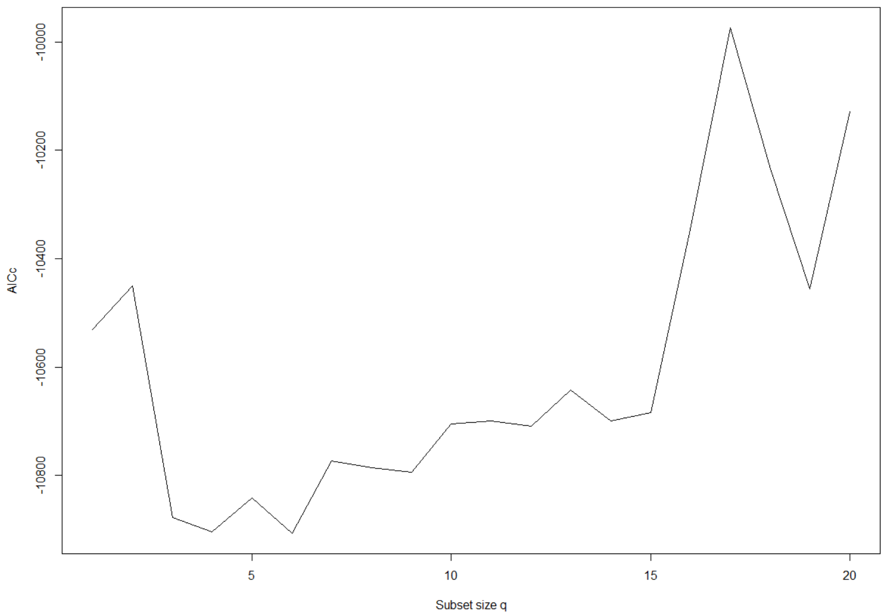


*Figure 6.7: The AICc score for optimal models of size q, ranging from 1-20.*

The selection process recognized that the optimal subset size was six explanatory variables with an AICc score of -10906. The aforementioned subset encompassed the following variables, smolt at 24 months lag, sea lice at lag three, consumption in the USA, fish loss at lag one, biomass at lag six, and temperature at lag three. These variables create the optimal model with regard to GAM forecasting. On the one hand, it is not possible to pass any inference from the results. On the other hand, it is interesting to point out that several variables that were deliberated to be of high significance in chapter 3 were also found in the subset. For instance, smolt release is conceivably the driving feature for future supply, seen as all future supply is predicated on how many smolt are released in the preceding years. Although causality is not possible to determine in this study, it is of high value to understand the variables, which, in conjunction, create an optimal model performance. Additionally, it is possible to break down the effect of each variable.

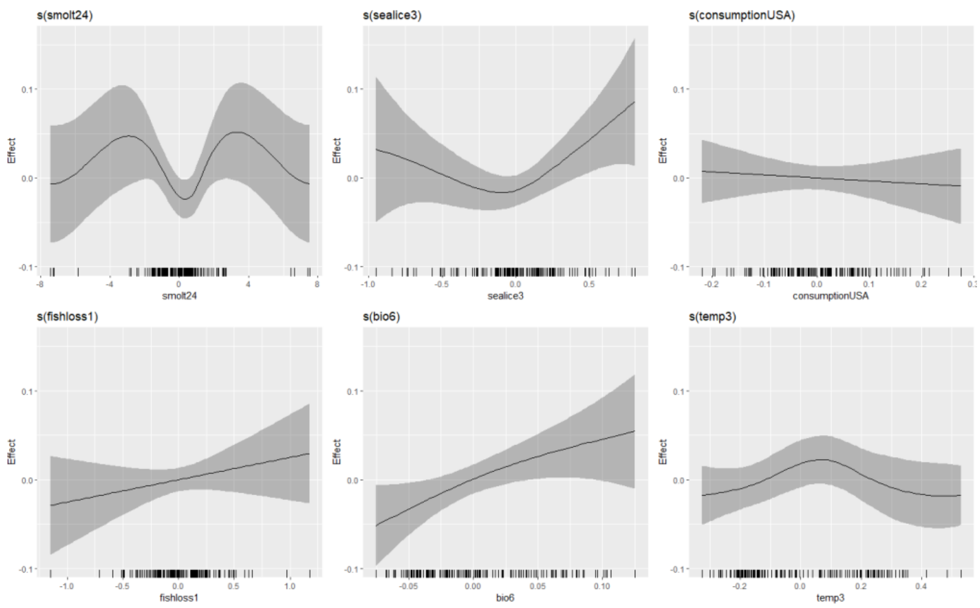


*Figure 6.8: The effect of each variable.*

As explained in section 4.2.2, GAM allows for non-linear terms. However, certain variables were shown to have a linear effect on the FPI, such as consumption, fish loss, and to an extent biomass. This is understandable, given that these variables are more stable throughout the year. The remaining features on the other hand, show nonlinearity. This is to be expected, given that smolt release and sea temperature has a seasonal cycle pattern, while sea lice can spread exponentially. That GAM allows for both nonlinear and linear fit is one of its major benefits when the real relationship is, in fact, non-linear. Furthermore, it is noteworthy that the optimal model fit would include consumption in the US rather than the EU, seen as the European Union is a substantially larger market. For future research, it would be interesting to create event studies that could isolate and study each feature separately.

The next step of the modelling was to fit data optimally. This was conducted in R with help from the package “mgcv”. In the same way as LASSO, the data was partitioned into the training set and test set. Within the training set, the nested cross-validations were performed for optimal results in tuning each spline.

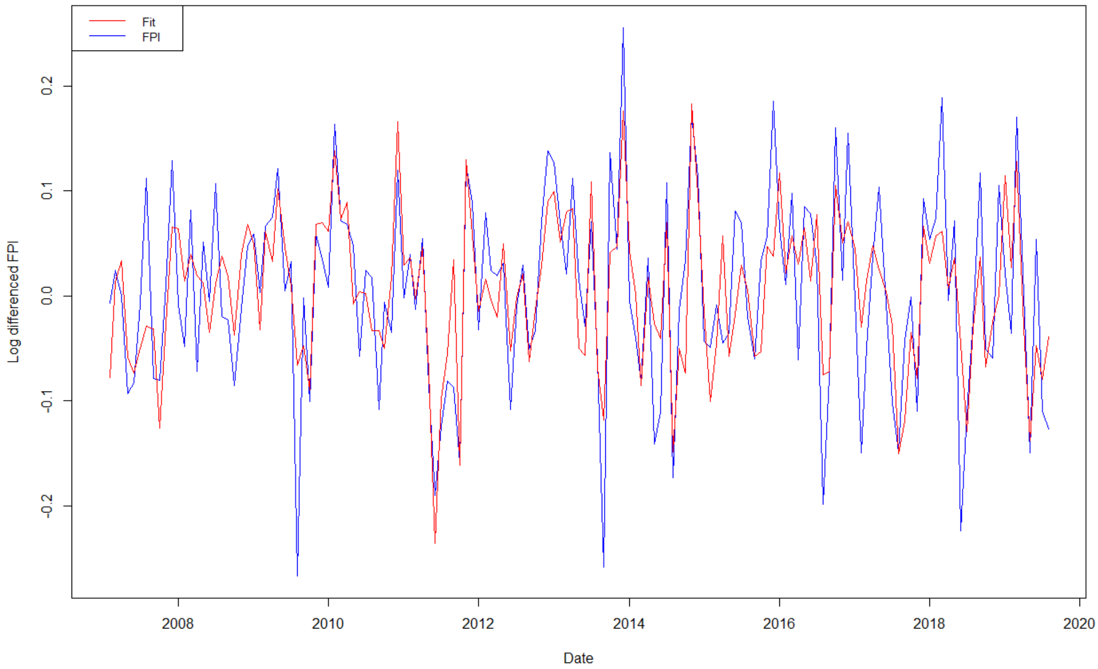


*Figure 6.9: GAM fit.*

From Figure 6.9, it is possible to see there is a close fit to the data throughout the timespan. However, to measure how close the fit is to FPI, we will measure the accuracy through MAE and RMSE.

*Table 6.14: GAM fit.*

| MAE | RMSE |
|---|---|
| 0.046 | 0.059 |

As displayed in Table 6.14, both measured in MAE and RMSE exhibit a relatively close fit. However, the real value is found in the results returned from out of sample prediction for 12-steps ahead. The data is all reverted back to the original format as done for the LASSO section 6.2.1.

*Table 6.15: GAM forecast 12-steps ahead.*

| Horizon | MAE | RMSE |
|---|---|---|
| 1 | 2.839 | 2.839 |
| 2 | 1.976 | 2.156 |
| 3 | 4.130 | 5.180 |
| 4 | 6.927 | 8.875 |
| 5 | 9.459 | 11.822 |
| 6 | 9.785 | 11.755 |
| 7 | 9.681 | 11.408 |
| 8 | 8.971 | 10.765 |
| 9 | 9.499 | 11.133 |
| 10 | 10.298 | 11.922 |
| 11 | 10.058 | 11.599 |
| 12 | 9.715 | 11.237 |

The results demonstrate inadequate predictions for from 5 months onwards measured in both MAE and RMSE, deviating only slightly after that from month to month. For example, the lowest error was observed in horizon two for all metrics at 1.976 for MAE and 2.156 for RMSE, while the highest was in horizon ten measured with an error of 10.298, 11.922 respectively.

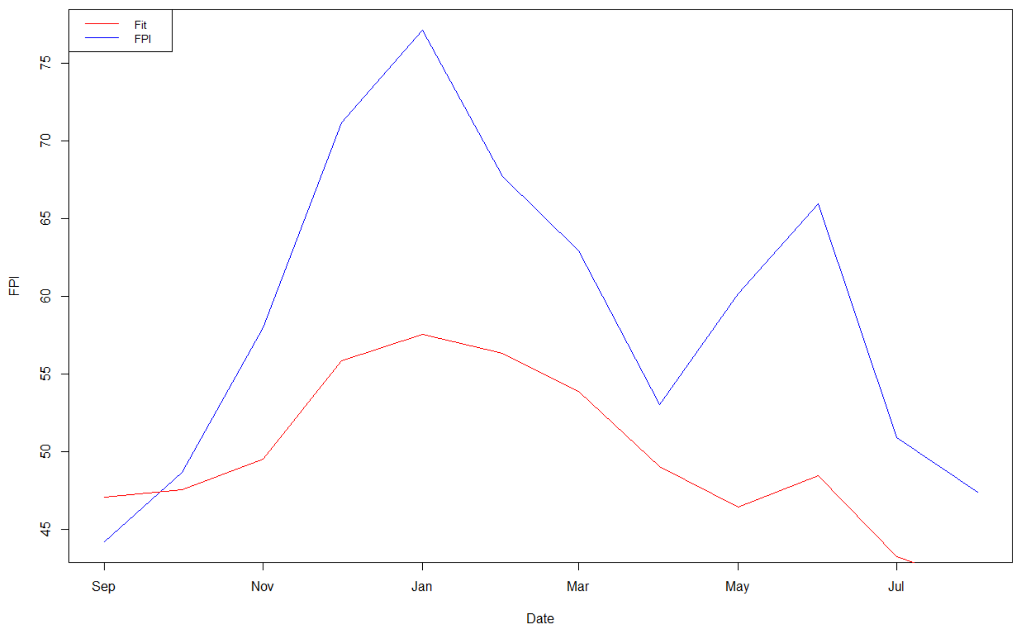


*Figure 6.10: Plot of the out of sample data and the forecast.*

Furthermore, the plot above illustrates that the model predicted the data exceptionally in regard to direction. The directional movements were accurately predicted 91.6% throughout the year horizon.

### 6.2.3 Multivariate Method Comparison

*Table 6.16: Multivariate model comparison.*

| | MAE | | RMSE | |
|---|---|---|---|---|
| | GAM | LASSO | GAM | LASSO |
| 1 month | 2.839 | 2.339 | 2.839 | 2.339 |
| 2 months | 1.976 | 2.452 | 2.156 | 2.454 |
| 3 months | 4.130 | 5.122 | 5.180 | 6.364 |
| 4 months | 6.927 | 8.901 | 8.875 | 11.522 |
| 5 months | 9.459 | 11.812 | 11.822 | 14.705 |
| 6 months | 9.785 | 11.958 | 11.755 | 14.389 |
| 7 months | 9.681 | 11.207 | 11.408 | 13.560 |
| 8 months | 8.971 | 10.164 | 10.765 | 12.725 |
| 9 months | 9.499 | 9.615 | 11.133 | 12.123 |
| 10 months | 10.298 | 9.828 | 11.922 | 12.085 |
| 11 months | 10.058 | 9.158 | 11.599 | 11.547 |
| 12 months | 9.715 | 8.846 | 11.237 | 11.165 |

From Table 6.16, it is possible to determine that GAM performed better than LASSO across mid-range horizons, while LASSO is superior in the first and last 2-3 months, measured in both MAE and RMSE. For instance, the average percentage difference in error over a 12-step horizon was 31% and 28% for MAE and RMSE, respectively. This underscores the superior performance from GAM in comparison to LASSO with regard to the magnitude of FPI forecasting over the entire horizon. The reason for this is difficult to pinpoint, however, one of the principal plausible causes could be the nonlinear nature from which the FPI derives. This is especially evident from Figure 6.7, which highlights the relationship between the features optimally suited for GAM forecasting. It showed how half of the predictors used in the analysis derived non-linear relationships when applied to the FPI. Although it is only evidence, it emphasizes that perhaps LASSO was inadequate in capturing substantial amounts of the real relationship given its parametric character. For example, LASSO selected several overlapping features with GAM, such as smolt release and sea temperature, both of which show non-linearity in relation to the FPI. This relationship would be forced into a linear one when modelling LASSO, essentially limiting their usage. In addition, GAM also predicted the directional movements better with 91% correct compared to LASSO with 50 %. The results illuminate the value in which industry players, especially salmon farming companies, can utilize multivariate methods in order to reduce the risk associated with a highly homogenous product. To illustrate, if a salmon farming company knew with a good deal of predictability the magnitude and directional movement of the FPI, they could reduce the harvest and/or tie-less up in the futures market until the price increases. Adversely, one can harvest more currently if the prices are predicted to drop and hedge risk by selling more through the futures market.

## 6.3 Univariate and Multivariate Method Comparison

We evaluate how these models perform and compare them against each other with the same metrics from earlier.

*Table 6.17: Forecast error for the lead univariate and multivariate methods.*

| | MAE | | | | RMSE | | | |
|---|---|---|---|---|---|---|---|---|
| | GAM | LASSO | Naïve | Arima | GAM | LASSO | Naïve | Arima |
| 1 month | 2.839 | 2.339 | 2.417 | 4.051 | 2.839 | 2.339 | 2.417 | 4.051 |
| 2 months | 1.976 | 2.452 | 4.174 | 7.035 | 2.156 | 2.454 | 4.529 | 7.641 |
| 3 months | 4.130 | 5.122 | 3.580 | 6.730 | 5.180 | 6.364 | 3.947 | 7.170 |
| 4 months | 6.927 | 8.901 | 4.336 | 5.312 | 8.875 | 11.522 | 4.754 | 6.232 |
| 5 months | 9.459 | 11.812 | 6.176 | 6.149 | 11.822 | 14.705 | 7.396 | 7.008 |
| 6 months | 9.785 | 11.958 | 8.272 | 7.030 | 11.755 | 14.389 | 10.209 | 7.919 |
| 7 months | 9.681 | 11.207 | 8.593 | 6.071 | 11.408 | 13.560 | 10.254 | 7.333 |
| 8 months | 8.971 | 10.164 | 7.864 | 6.317 | 10.765 | 12.725 | 9.641 | 7.424 |
| 9 months | 9.499 | 9.615 | 7.796 | 7.452 | 11.133 | 12.123 | 9.406 | 8.909 |
| 10 months | 10.298 | 9.828 | 7.034 | 7.170 | 11.922 | 12.085 | 8.923 | 8.578 |
| 11 months | 10.058 | 9.158 | 7.086 | 6.910 | 11.599 | 11.547 | 8.812 | 8.281 |
| 12 months | 9.715 | 8.846 | 6.975 | 6.862 | 11.237 | 11.165 | 8.598 | 8.137 |

From Table 6.17, it is possible to determine that the univariate models outperformed the multivariate methods across the entire time horizon except for the first two months in accuracy. Conversely, LASSO outperformed GAM in the first month, while GAM was superior in month two. GAM performed best with regard to directional movement, followed by naïve, ARIMA, and LASSO. ARIMA showed superior results of the two univariate methods in medium to long-range (i.e., 5-12 months) except for month 10 if measured in MAE. However, the difference here is minor. Overall, ARIMA is preferable, followed by naïve, GAM, and LASSO.

# 7 Discussion

This thesis has focused on exploring several previously utilized univariate models in comparison with two regression methods with regard to the FPI. The objective was to explore how the inclusion of univariate and multivariate methods could create solid forecast for the salmon price over a 12-month period. To begin with, we created two simple forecasting benchmark models, which were the naïve and rwdrift method in combination with the STL decomposition to seasonally adjust the benchmark models. For the additional univariate forecasting methods, we chose to include the ARIMA and ETS model, as both forecasting methods are suitable for forecasting time series with both trend and seasonality. Lastly, for the multivariate forecasting methods, we included LASSO and GAM. All models created a forecast for the time interval 1 to 12 months ahead.

## 7.1 General Findings

From the analysis, it is shown that none of the methods performed particularly well, with the exception of the first two months. However, how poorly they performed is a matter of discussion, with the error rate mostly varying between 7-11 for the multivariate methods and 4-9 for the univariate methods measured in MAE. This can be seen in comparison to the FPI varying between NOK44 and NOK77. In other words, the univariate models performed substantially better across the timeframe of 3-12 months. There can be a plethora of reasons for why the univariate models performed better or the multivariate performed worse.

Firstly, the results indicate that multivariate models in both parametric LASSO and partially non-parametric GAM are unable to fully capture the data adequately. In different terms, more complex models do not necessarily entail increased performance, as evident from the results in Table 6.17. However, a different explanation could be found in the sample size. Section 4.3 discussed how a sufficient sample size was paramount to achieve accurate predictions. This might explain why LASSO did poorly across much of the timeframe, given that all explanatory variables, including the lags, were included when performing the LASSO. In other words, LASSO was modelled with a restriction on the parameter space, strongly suggesting that additional observations could improve accuracy. However, seen as most of the data could be found at a monthly frequency at the lowest, it would be impossible to produce more data for the different time series. A different approach could be to limit the number of lags introduced. However, this could entail leaving out important variables. Moreover, GAM was able to accurately predict the directional movements 91.6% of the time compared to 50% for LASSO, further indicating GAM was able to fit the data better.

Furthermore, if the real relationship in the data were actually represented by fewer parameters, it would have negative ramifications in the prediction using LASSO. However, it is impossible to determine the real relationship to non-simulated data. Another possible reason why the univariate methods outperformed LASSO and GAM is the fact that most of the data has high variance, making it harder for more complex multivariate models to sufficiently capture the data. In other words, it would seem the simplicity of ARIMA, and especially naïve is of value rather than a cause of detriment. As elaborated earlier, the ARIMA model is more suitable for smaller time series with a short forecast horizon. The ARIMA model is a combination of an Autoregressive model (AR) and a Moving average model (MA), which are integrated by differencing, which seems to suit the time series well. The autoregressive part of the model is remarkably flexible for different time series, and the moving average part of the model uses past forecast errors to optimize the forecast, which our time series benefits from. Lastly, the seasonal variation and trend component from the dataset is also coped with by differencing the time series and making it stationary. Since the ARIMA model performs just as well as or better than the multivariate forecasting methods, it is reasonable to assume that the historical data of the salmon price explains a lot of the variation in the price, and therefore has high predictive power.

Moreover, the results also indicate that GAM outperformed LASSO given the non-linear character as seen from half the GAM selected predictors, while the rest exhibiting linearity in the relationship with the response (see Figure 6.7). In other words, LASSO would struggle to fit much of the data, while GAM would allow for nonlinear fit, resulting in GAM performing better. Unlike the univariate models, GAM and LASSO did provide valuable insight into which features provide stronger predictive power. For instance, both had several overlapping variables from the selection process. These were smolt release at lag 24 and sea temperature at lag three, which raises the question of why these two features were included in both LASSO and GAM. There could be a causal relationship here, which, if understood more closely, could provide insight into how salmon farming companies could, for example, structure their smolt release timeline and in which waters the facilitates are set. Despite the indication of strong predictive power, it is not possible to infer any causality from the results. This is outside the scope of this paper. However, the results strongly suggest further research into especially these two explanatory variables.

## 7.2 Scenario Analysis

Arima was shown to be the most optimal method. However, how beneficial would our forecast be for a single salmon farming company? We will use a hypothetical company called Dagslaks with a Norwegian market share of 20% to illustrate. ARIMA was in the short term and medium-term able to predict the first 5 months satisfactory. This could be of substantial value for Dagslaks, given knowing the directional movements and having relatively good accuracy would help with implementing strategic choices. For instance, in our case, the FPI increases from about NOK44 to NOK77 over the first 5 months. Where all of those months saw an increase in price, knowing this with some certainty, Dagslaks could reduce harvesting and then slowly increase harvest output as the price rises. Doing this would entail selling the salmon at a substantially higher price. However, the difficult question is how much Dagslaks should reduce production and how fast it should increase it. A potential option could be to reduce harvest by 30% (the equivalent of 3216 tonnes for Dagslaks) the first month, before increasing in increments of 5% of total production and lastly selling all excess salmon kept for the highest price point at month 5. Mirroring 20% of real actual production, we get that Dagslaks could have increased revenue by 235.6 million over a 5-month period (see appendix A.3 Scenario Analysis). On the one hand, this analysis is a rough estimate and takes several liberties in its calculation. On the other hand, it highlights what could be gained by salmon farming companies if they were to create superior forecasts. However, this isolated analysis does not account for negative externalities. For example, to what extent holding back harvest would negatively impact the price or considering the operational side of harvest all held off salmon the last month in our case.

In the long-term, ARIMA forecast accuracy was stable, but the directional movements were poor. However, given that the forecast had been right in predicting a price reduction after the medium range, then DagsLaks could have scaled back planned investment in new facilities, reduce the number of smolt released, and implemented a hiring stop. On a different note, making long-term decisions on faulty assumptions of the scale of the movements would impact the bottom line by making consequential choices when they were, in fact, unnecessary. In essence, these are just some of many examples of how our forecast, and good forecasts can be of great value to salmon farming companies.

## 7.3 Forecast Model Error

As elaborated above, all forecast methods have some errors compared to the test set with the actual observations. We want to look into the data set to determine if any of the forecast errors can be explained by the data. Figure 7.1 displays the actual observations of the salmon price from 2019, which includes our forecast horizon.

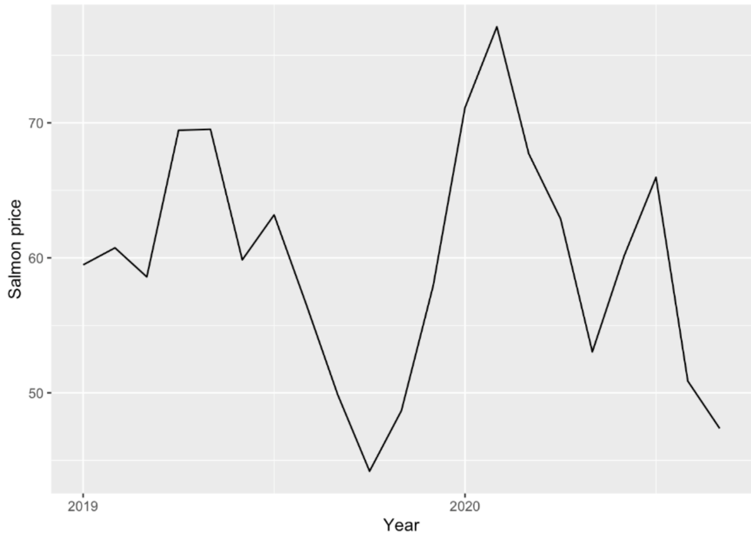


*Figure 7.1: FPI 2019-2020.*

From Figure 7.1, it is possible to observe a sharp price fall from February 2020. At this time, the world experienced the outbreak of the Covid-19 virus, as elaborated in section 5.1. This virus outbreak led to reduced demands, thus reduced export and corresponding price fall. Further, in section 5.1, we discussed the import ban of salmon in Russia from 2014 was bypassed by importing through Belarus. However, this exception had been discussed to be shut down from the second half of 2019 and was finally closed down in January 2020, which is where the salmon price had a sharp fall in the time series. It is, therefore, reasonable to assume that this pandemic and political noise has affected the time series and the salmon price. The forecasting models we built seems to struggle to pick up such political and global noise, as it is not reflected on any of the variables used in this paper. Given that this assumption is correct, it is possible that the forecast models would have performed better in another time interval.

## 7.4 Limitations

### 7.4.1 Sparse Data

Perhaps the largest limitation of the study is the amount of scarce data. In the univariate methods, 176 observations are used, while the multivariate methods used 151 observations. The sparse data becomes more evident when one takes into account that a large number of parameters are added to the modelling. This is especially the case for multivariate forecasting, where lagged features remove the number of observations equal to their lag. Furthermore, the nested cross-validation is restricted, given that the training set should ideally start with more than half of the observations before rolling it towards the end of the training data. As a consequence, a limited amount of k-fold along time series is possible. In essence, the extent of the time series data might be inadequate to capture the fit in an optimal fashion.

### 7.4.2 Variation

In section 2.2.1.2 and 5.3, it was described and shown how especially FPI but also most of the explanatory variables exhibited strong and unequal levels of variation throughout the time series (i.e., highly volatile). This was, to some extent, accounted for using log transformation. However, what is not possible to incorporate is the random variation in the data caused by being subject to external factors such as weather, air quality, etc. Therefore, it is difficult to discount that despite the transformation, randomness will still have a lasting impact on the forecasts. This is especially the case with regard to multivariate forecasting, as briefly mentioned in section 7.1. The reason lies in the fact that multivariate models include 55 independent variables.

### 7.4.3 Inaccessible Variables

In chapter 3, it was discussed which variables are most important when forecasting the FPI. However, additional features could be introduced for future analysis, which were unattainable at the time of writing. For instance, many potential variables were either not accessible or would have required an extraneous amount of manual data gathering or estimations, which itself would have taken months to complete. Examples of these include predictors connected to production in Chile, the second-largest producer after Norway. Data pertaining to Chilean production was often scattered, not recorded, or often held as company secrets, therefore unattainable for public usage. Another example would be to break down biomass based generational make-up of salmon population to provide an improved picture of short-term future supply. However, such data was not to be found. Therefore, futures analysis could be improved by encompassing features currently unattainable, such as the ones mentioned above.

## 7.5 Possible Improvement

### 7.5.1 Sample

A possible improvement for future analysis could be to introduce additional data. In section 7.1, it was mentioned how it was not possible to add new data to the time series. However, this was only the case because of the number of variables included. One could potentially reduce the variables to a subset from which all the selected features had higher frequency data, the one alternative being weekly given the FPI. Although this could increase the amount of data, it would be at the expense of discarding predictors, which could have strong predictive power. For example, if the analysis were to include weekly observations, only futures prices, exchange rates, and index/prices for alternative proteins would be included. Following this example, the number of observations would increase to 764. The same would also apply to univariate models. Nevertheless, it would not guarantee improved accuracy, and it would defeat part of the aim of the thesis to include a diverse set of features to see which has strong predictive power.

### 7.5.2 Consequences of the Delimitations

In chapter 4, we elaborated on the characteristics of our forecasting methods and why we believed these methods would suit the dataset. However, from the forecast results and accuracy measures from the different forecasting methods, we are able to see that all models have some errors and are not capable of fully learning the seasonal- and trend patterns. This indicates that the dataset may not suit these forecasting methods as well as first thought. As mentioned in section 2.2.1.3, we chose to restrict this thesis to compare proven forecasting models from previous literature against unexplored forecasting methods. This thesis has therefore solely focused on the four forecasting methods, ETS, ARIMA, GAM, and LASSO, in addition to the benchmark forecast models, naïve and rwdrift with STL decomposition. As a consequence of our delimitation, it is possible that we have overlooked some forecasting methods which could have produced more accurate forecasts. An example would be a neural network, which is known for great computational power and advanced algorithms. However, one of the advantages of this method is that it is great at predicting when there is a large dataset provided, which is not the case in our study. Secondly, a well-known disadvantage of the neural network model is the “black box” phenomenon, which means that the model has poor human interpretability of what the prediction is based on. And since we want to interpret the model and the variable importance, this model was not chosen. Methods like decision trees and random forest were also considered, but these models suit classification problems better than regression problems; hence they were rejected.

As we have observed in the findings of this paper, the ARIMA model has performed quite well as a univariate model. The ARIMA model could also be used as a multivariate model in R. Implementing explanatory variables in the ARIMA model could increase the forecast precision of an already well-performing model. However, we have chosen to include the ARIMA model only as a univariate model and not multivariate, since we only want to explore the forecast methods chosen initially. The multivariate forecast method should, however, be included in future research.

Based on the forecasting methods we chose to explore in this study and their strengths, we have chosen to restrict the forecasting horizon to a short to medium forecast, from 1 to 12 months ahead, thus choosing to not forecast for longer time periods, such as 2 and 3 years ahead forecast. Hence, a possible improvement could be to forecast for longer periods with either the forecasting methods chosen in this study or other methods that were not covered.

A large part of multivariate forecasting was to decide on variables deemed to impact the FPI (see chapter 3). However, several potential variables were left out during this process. Given that they were thought to be of less importance and/or to avoid multicollinearity. The predictive power of the model is not reduced by collinearity, however, it will have a detrimental impact on the predictor interpretability. Hence, this thesis tried to leave out strong collinear variables, such as biomass measured in individuals versus the one used, which is biomass measured in tonnes, sea temperature at different depths, etc. Consequently, future research could try to incorporate or replace variables included in this analysis to study for disparate effect.

# 8 Conclusion

The growth of the salmon industry has made it an integral part of the Norwegian economy. As a consequence, salmon farmers and participants along the value chain can gain considerable insight from forecasting the price of salmon and a greater understanding of the fundamental features with the strongest predictive power. Therefore, we set out to answer the following research question: *Can the implementation of univariate and multivariate time series forecast methods create solid forecasts for the price of salmon 12 months ahead?*

Firstly, previously tried univariate methods were used with updated price information. These included autoregressive integrated moving average (ARIMA) and exponential smoothing (ETS) in addition to simple benchmark methods, naïve and rwdrift with STL decomposition. Secondly, two multivariate methods were applied, least absolute shrinkage and selection operator (LASSO) and generalized additive model (GAM). In this study, the Fish Pool Index (FPI) was chosen to represent the spot price of Atlantic salmon. The additional features included in the multivariate models were incorporated based on a thorough independent analysis, the Salmon Farming Industry Handbook, and consultation from Kontali Analyse AS.

The general findings show that overall, both univariate and multivariate methods do not produce accurate forecasts over a 12-step horizon. Furthermore, the naïve benchmark and ARIMA outperform the multivariate models, with the exception of the first 2 months. Where naïve does better short-term, while ARIMA is superior medium to long-term. Among the multivariate methods, GAM showed smaller errors when compared to LASSO. This might be caused by the intrinsic non-linear relationship between several of the predictors and the FPI, which LASSO is unable to capture. Nevertheless, the results are a strong indicator that simpler univariate models are preferred with regard to the FPI. However, GAM and LASSO provided an improved understanding of the predictors most important to forecast the FPI for the respective methods. For instance, both models selected smolt release and temperature at lag 24 and three, respectively. Although it is not an inference of causal relationships, it does provide a foundation for which additional research can be conducted. In essence, neither univariate nor multivariate methods provide adequate forecasts to solely base important strategic decisions. Consequently, industry participants would need to supplement with additional insight and analysis before making important long-term decisions.

The reasons for why the univariate models outperformed the multivariate are difficult to determine. However, several hypotheses are possible. Firstly, the scarce amount of data would seem to have a stronger adverse effect on GAM and LASSO compared to naïve and ARIMA, given the high ratio of parameters to observations. Furthermore, there seems to be a substantial amount of random variation in many of the variables, which is thought to have an accumulated negative effect on the multivariate models. Lastly, it is important to emphasize that many avenues of salmon forecasting remain unexplored, either with regard to the methods presented here or completely different approaches discussed earlier. It is paramount to continue researching in order to contribute invaluable insight in which salmon farming companies and other industry participants can base decision making.

# 9 References


Adebiyi, A. A., Adewumi, A. O., & Ayo, C. K. (2014, March 05). *Comparison of ARIMA and Artificial Neural Networks Models for Stock Price Prediction*. Retrieved from Hindawi: https://www.hindawi.com/journals/jam/2014/614342/

Ankamah-Yeboah, I., Nielsen, M., & Nielsen, R. (2017). Price formation of the salmon aquaculture futures market. *Aquaculture Economics & Management* , 376-399.

Asche, F., Misund, B., & Oglend, A. (2016). The spot-forward relationship in the Atlantic salmon market. *Aquaculture Economics and Management*, 222-234.

Bloodworth, J., Baptie, M., Preedy, K., & Best, J. (2019). *Negative effects of the sea lice therapeutant emamectin benzoate at low concentrations on benthic communities around Scottish fish farms*. Retrieved September 2020, from http://www.sciencedirect.com/science/article/pii/S0048969719309428

Bloznelis, D. (2017). Short-term salmon price forecasting. *Journal of Forecasting*.

Burnham, K. P., & Anderson, D. R. (2002). *Model Selection and Multimodel Inference: A Practical Information-Theoretical Approach, Second Edition* . New York: Springer.

Cancer Council. (2018). *Red meat, processed meat and cancer*. Retrieved from Cancer council: https://www.cancercouncil.com.au/1in3cancers/lifestyle-choices-and-cancer/red-meat-processed-meat-and-cancer/

Central Bank of Norway. (2020, August). *Exchange Rates*. Retrieved from Central Bank of Norway: https://www.norges-bank.no/en/topics/Statistics/exchange_rates/?tab=currency&id=EUR

Dickey, D., & Fuller, W. A. (1995). Lag order and critical values of the augmented Dickey-Fuller test. *Journal of Business & Economic Statistics* , 277-280.

Directorate of Fisheries. (2020, September). *Nøkkeltall for norsk havbruksnæring*. Retrieved from Directorate of Fisheries: https://www.fiskeridir.no/Akvakultur/Tall-og-analyse/Statistiske-publikasjoner/Noekkeltall-for-norsk-havbruksnaering

Falconer, L., Hjøllo, S. S., Telfer, T., Mcadam, B., Hermansen, Ø., & Ytteborg, E. (2020). *The importance of calibrating climate change projections to local conditions at aquaculture sites*. Retrieved September 2020, from http://www.sciencedirect.com/science/article/pii/S0044848619316199

Federal Reserve Bank of Saint Louis. (2020, September). *Global price of lamb*. Retrieved from Federal Reserve Bank of Saint Louis: https://fred.stlouisfed.org/series/PLAMBUSDM

Fish Pool. (2020). *Fish Pool Index*. Retrieved from Fish Pool: https://fishpool.eu/price-information/spot-prices/

Fish Pool. (2020, August 25). *Forward price history*. Retrieved from Fish Pool: https://fishpool.eu/price-information/forward-prices-3/forward-closing-prices-history/

Fish Pool. (2020, August 2020 29). *Forward Prices*. Retrieved from Fish Pool: https://fishpool.eu/price-information/forward-prices-3/

Fish Pool. (2020, August 25). *Price history*. Retrieved from Fish Pool: https://fishpool.eu/price-information/spot-prices/history/

Glantz, M., & Mun, J. (2011). *Extrapolation*. Retrieved from Sciencedirect: https://www.sciencedirect.com/topics/social-sciences/extrapolation

Global Salmon Initiative. (2017). *About salmon farming*. Retrieved from Global Salmon Initiative: https://globalsalmoninitiative.org/en/about-salmon-farming/

Gu, G., & Anderson, J. (1995). Deseasonalized State-Space Time Series Forecasting with Application to the US Salmon Market. *Marine Resource Economics*, 171-185.

Guttormsen, A. G. (1999). Forecasting weekly salmon prices: risk management in fish farming. *Aquaculture Economics & Management* , 159-166.

Harrel, F. E., Lee, K. L., Califf, R. M., Pryor, D. B., & Rosati, R. A. (1994). Regression Modelling Strategies for Improved Prognostic Prediction. *Statistics of Medicine* , 143-152.

Holland, J. (2020, November 4). *Lower salmon prices, increased costs impact Norway Royal Salmon's Q3 earnings*. Retrieved from Sea Food Source: https://www.seafoodsource.com/news/business-finance/lower-salmon-prices-increased-costs-impact-norway-royal-salmon-s-q3-earnings

Hyndman, R. J., & Athanasopoulos, G. (2018). *Forecasting: Principles and Practice.* Monash University, Australia: Otexts.

Hyndman, R. J., & Kostenko, A. V. (2007). Minimum Sample Size Requirements for Seasonal Forecasting Models. *Foresight*.

International Salmon Farmers Association. (2018, June). *Salmon Farming: Sustaining communities and feeding the world.* Retrieved from Sjømatnorge: https://sjomatnorge.no/wp-content/uploads/2018/06/ISFA-Report-2018-FINAL-FOR-WEB.pdf

Intrafish. (2019, November 4). *Is Chile's salmon industry cursed?* Retrieved from Intrafish: https://www.intrafish.com/commentary/is-chiles-salmon-industry-cursed-/2-1-700462

James, G., Witten, D., Hastie, T., & Tibshirani. (2017). *An Introduction to Statisical Learning* . Los Angeles : Springer.

Kwiakowski, D., Phillips, C., Schmidt, P., & Shin, Y. S. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics* , 159-178.

Lin, B.-H., Herrman, M., Lin, T.-Y., & Mittelhammer, R. C. (1989). Forecasting the price of farmed Atlantic salmon: An integrated econometric and time series approach . *Agribusiness*, 477-488.

Lusedata. (2020, September). *Lusedata in Excel*. Retrieved from Lusedata: https://lusedata.no/statistikk/excel/

McQuarrie, A. D., & Tsai, C.-L. (1998). *Regression and Time Sereis Model Selection.* Singapore: World Scientific Publishining .

MOWI. (2020). *Salmon Farming Industry Handbook*. Retrieved September 2020, from https://mowi.com/it/wp-content/uploads/sites/16/2020/06/Mowi-Salmon-Farming-Industry-Handbook-2020.pdf

Mundi, Index. (2020, September). *Poultry Daily Price*. Retrieved from Index Mundi: https://www.indexmundi.com/commodities/?commodity=chicken

Norwegian Directorate of Fisheries. (2020, August). *Aquaculture*. Retrieved from Norwegian Directorate of Fisheries: https://www.fiskeridir.no/Akvakultur

Norwegian Directorate of Fisheries. (2020, September 5). *Matfisk produksjon av laks, regnbueørret og ørret.* Retrieved from fiskeridir: https://www.fiskeridir.no/Akvakultur/Tall-og-analyse/Akvakulturstatistikk-tidsserier/Laks-regnbueoerret-og-oerret/Matfiskproduksjon

Norwegian Seafood Council. (2020, January 7). *Norwegian seafood exports top NOK 107 billion in 2019*. Retrieved from Seafood: https://en.seafood.no/news-and-media/news-archive/norwegian-seafood-exports-top-nok-107-billion-in-2019/

Oglend, A. (2013). Recent trends in salmon price volatility. *Aquaculture Economics & Management* .

Oglend, A., & Sikveland, M. (2008). The Behavior of Salmon Price Volatility . *Marine Resource Economics* , 507-526.

Penn state. (2020). *8.2 Cross Correlation Functions and Lagged Regressions*. Retrieved from Penn State: Eberly College of Science: https://online.stat.psu.edu/stat510/lesson/8/8.2

Quandl. (2020, September). *Beef; Australian and New Zealand*. Retrieved from Quandl: https://www.quandl.com/data/ODA/PBEEF_USD-Beef-Australian-and-New-Zealand-85-lean-fores-CIF-U-S-import-price-US-cents-per-pound

Salmon Business. (2019, December 30). *Russia stops import of Norwegian salmon and trout via Belarus*. Retrieved from Salmon Business: https://salmonbusiness.com/russia-stops-import-of-norwegian-salmon-and-trout-via-belarus/

Sestelo, M., Villanueva, N. M., Meira-Machado, L., & Roca, J. (2016). FWDselect: An R Package for Variable Selection in Regression Models. *The R Journal*, 132-148.

Shumway, R., & Stoffer, D. (2016). *Time Series Analysis and Its Applications With R Examples* (Vol. 4.). Davis, California: Springer.

Solstad, T. L. (n.d.). *4 reasons why the salmon market will experience a significant growth*. Retrieved from Qurillaqua: https://www.qrillaqua.com/blog-and-news/4-reasons-why-the-salmon-market-will-experience-a-significant-growth

Tibshirani, R. (1996). Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society* , 267-288.

Vukina, T., & Anderson, J. L. (1994). Price forecasting with state-space models of nonstationary time series: case of the Japanese salmon market. *Computers & Mathematics with Applications*.

Worldometer. (2020, November). *World population*. Retrieved from Worldometer: https://www.worldometers.info/world-population/

# 10 Appendix

## A.1 Data Sources

**Kontali Analyse AS**

Kontali Analyse AS is a private company with a strong focus and expertise in the aquaculture and seafood industry. They generously provided data not accessible publicly; these include harvest volume (tonnes) excluding Norway, price of trout, consumption in EU, USA, Japan, and one for the remaining market. Our contact person in Kontali Analyse AS was Kunt Henrik Rolland.

**Fishpool**

Fishpool is an international salmon exchange based in Norway and owned by Oslo Stock Exchange. The company operates as a marketplace for the selling and purchase of salmon contracts. Used here was the average monthly spot price for 3-6kg head-on gutted. Historically the FPI has consisted of a weighted average from NASDAQ prices FCA Oslo and SSB custom export statistics. Throughout the years, it has also incorporated farmers selling price FOB and Fish Pool European buyers index to a varying extent. The weight distribution continuously changes, however, the NASDAQ salmon price index is the main weight. The forward prices are also pulled from here.

Accessible at: (Price history, 2020) & (Forward price history, 2020)

**Directorate of Fisheries**

The Directorate of Fisheries' mandate is to facilitate a sustainable and profitable fish industry. Our data on biomass (tonnes), smolt release (number of individuals), and harvest quantity monthly (number of individuals) basis was collected from their database. All the data is reported at the end of each month. Harvest is recorded by the number of individuals and could be converted into tonnes if multiplied by the average harvest weight of 4.5, followed by a division of 1000.

Accessible at: (Aquaculture, 2020)

**Quandl**

Quandl is a company which provides data for business to make improved decisions. The data gathered was given as a beef value index of New Zealand/Australian exports.

Accessible at: (Beef; Australian and New Zealand, 2020)

**Index Mundi**

Index Mundi is a data portal that structures raw data into useful, structured data, making it accessible to anyone requiring information. Price on whole chickens was given in dollars/kg.
Accessible at: (Poultry Daily Price, 2020)

**Saint Louis Federal Reserve**

The Saint Louis Federal Reserve Bank research division is a high-quality provider of top research on macroeconomics, applied microeconomics, and finance. Obtained here was the global price of lamb (largest exporter) in cents/pound monthly average, which was subsequently altered to dollars/kg.
Accessible at: (Global price of lamb, 2020)

**Norges-Bank**

The Norwegian Central Bank stabilizes the Norwegian economy through the highest level of economic decision making. They also provide statistics about exchange rates (NOK/EUR) and NOK/USD.
Accessible at: (Exchange Rates, 2020)

**Lusedata**

Lusedata is a service that provides valuable information to the aquaculture industry service. Data on sea temperature (Celsius) averaged across all Norwegian regions, average sea lice occurrences per salmonid, and sea lice treatment in percent of the total amount of salmonids.
Accesible at: (Lusedata, 2020)

## A.2 Descriptive Statistics

| Feature Information | Descriptive Statistics | | | | Root Tests | |
|---|---|---|---|---|---|---|
| *Variables* | *Observations* | *Std. deviation* | *Mean* | *FPI Correlation* | *KPSS* | *ADF* |
| Fish Pool Index (NOK/kg) | 188 | 14.980 | 41.08 | 1.000 | 0.319 | -1.679 |
| Biomass (Tonnes, Norway) | 188 | 141839.3 | 600501.1 | 0.618 | 0.418 | -1.558 |
| Smolt release (#Individual Salmon, Norway) | 188 | 17915830 | 21325150 | 0.087 | 0.038 | -7.296 |
| Sea Temperature (Celcius, Norway) | 188 | 3.127 | 8.81 | -0.066 | 0.014 | -10.086 |
| Bovine Index (Index) | 188 | 38.9 | 167.06 | 0.554 | 0.409 | -1.503 |
| Poultry Price (USD/kg) | 188 | 0.258 | 1.936 | 0.304 | 0.388 | -2.253 |
| Lamb Pice (USD/kg) | 188 | 0.584 | 2.88 | -0.710 | 0.261 | -2.186 |
| NOK/EUR | 188 | 0.865 | 8.59 | 0.750 | 0.655 | 0.080 |
| NOK/USD | 188 | 1.282 | 6.92 | 0.803 | 0.628 | -0.724 |
| USD/EUR | 188 | 0.08 | 0.799 | 0.708 | 0.378 | -1.952 |
| Harvest (#Individuals, Norway) | 188 | 5276657 | 17401450 | 0.561 | 0.237 | -1.982 |
| Fish Loss (#Individual Salmon, Norway) | 188 | 1403763 | 3664938 | 0.387 | 0.060 | -3.927 |
| Sea Lice Treatments (Share of Salmon) | 188 | 0.062 | 0.109 | -0.252 | 0.202 | -4.012 |
| Sea Lice (Average amount/Salmon) | 188 | 0.546 | 1.107 | -0.031 | 0.284 | -3.891 |
| Forward Prices (NOK/kg) | 188 | 14.95 | 42.49 | 0.726 | 0.215 | -1.600 |
| Price of Trout (NOK/kg) | 188 | 13.66 | 41.35 | 0.944 | 0.183 | -1.747 |
| Harvest (Tonnes, ex. Norway) | 188 | 20163.14 | 72689.00 | 0.670 | 0.400 | -0.881 |
| Consumption EU (Tonnes) | 188 | 16488.12 | 74117.53 | 0.529 | 0.117 | -2.009 |
| Consumption USA (Tonnes) | 188 | 7753.65 | 30426.56 | 0.747 | 0.709 | 0.639 |
| Consumption Japan (Tonnes) | 188 | 1115.666 | 4564.86 | 0.305 | 0.255 | -2.939 |
| Consumption Other (Tonnes) | 188 | 17351.97 | 48438.93 | 0.618 | 0.364 | -1.593 |

*ADF*: $\tau_\alpha = 5\% < -3.45$; $\tau_\alpha = 1\% < -4.04$

*KPSS*: $\tau_\alpha = 10\% > 0.119$; $\tau_\alpha = 5\% > 0.146$; $\tau_\alpha = 1\% > 0.216$

## A.3 Scenario Analysis

| | | Without reduction | | With reduction | | | |
|---|---|---|---|---|---|---|---|
| Month | Price/kg | Harvest Tonnes | Revenue NOK | Harvest Tonnes | Delayed Harvest This Month | Total Delayed Harvest | Revenue NOK |
| September | kr 44.20 | 10718 | 473 735 600.00 | 7502.6 | 3 215 | 3 215 | 331 614 920.00 |
| October | kr 48.68 | 10846 | 527 983 280.00 | 8134.5 | 2 712 | 5 927 | 395 987 460.00 |
| November | kr 57.94 | 11058 | 640 700 520.00 | 8846.4 | 2 212 | 8 139 | 512 560 416.00 |
| December | kr 71.13 | 11496 | 817 710 480.00 | 9771.6 | 1 724 | 9 863 | 695 053 908.00 |
| January | kr 77.11 | 10876 | 838 648 360.00 | 20738.9 | 0 | 0 | 1 599 176 579.00 |
| **Total** | | | **3 298 778 240.00** | | | | **3 534 393 283.00** |
| **Revnue Gain** | | | | | | | **235 615 043.00** |