Master's thesis

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

Martin Aas Granviken
Jan Fredrik Herud
Jørgen Valseth

Price and currency hedging strategies for Norwegian salmon producers

A GARCH modeling approach

Master's thesis in Industrial Economics and Technology Management
Supervisor: Maria Lavrutich
June 2020

NTNU
Norwegian University of Science and Technology

Martin Aas Granviken
Jan Fredrik Herud
Jørgen Valseth

Price and currency hedging strategies for Norwegian salmon producers

A GARCH modeling approach

Master's thesis in Industrial Economics and Technology Management
Supervisor: Maria Lavrutich
June 2020

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

NTNU
Norwegian University of
Science and Technology

Preface

This master thesis is conducted as part of achieving the degree Master of Science in Industrial Economics and Technology Management at the Norwegian University of Science and Technology (NTNU), Trondheim. The specialization is within Financial Engineering at the Department of Industrial Economics and Technology Management. The thesis is a result of independent work by Martin Granviken, Jan Fredrik Herud and Jørgen Valseth, and was undertaken from January to June 2021. The motivation behind the thesis is based on academic interest and the lack of scientific literature on hedging other financial risks than price risk in the aquaculture industry.

We want to express our profound gratitude to our supervisor Associate Professor Maria Lavrutich for valuable guidance and feedback, and stimulating discussions throughout the semester. We also want to thank our friends and family for their continuous support.

Trondheim, June 2021

Martin Aas Granviken
Jan Fredrik Herud
Jørgen Valseth

Abstract

With limited growth potential, combined with the inherently risky nature of the business, the need for efficient financial risk management in the Norwegian salmon farming industry is of growing importance. In this thesis, we address two of the most important financial risks salmon producers must manage to stay competitive and profitable, price risk and currency risk. Hence, we investigate a joint salmon price and currency hedging problem for a Norwegian salmon producer exporting to the EU. First, we examine the performance of dynamic strategies modeled with a state-of-the-art GARCH model compared to the traditional naïve hedge, and find that dynamic strategies are valuable for salmon producers, especially in terms of return. Second, we examine how a multi-product hedging framework that takes dependencies between salmon price and currency into account, perform in comparison to hedging them separately. We find that taking the dependencies into account is beneficial for hedging performance. Third, we introduce a novel "threshold strategy" that utilizes volatility clustering effects for hedging purposes. Our results show that this strategy outperforms the more traditional approaches in terms of returns and performs similarly in terms of risk reduction.

Keywords: Aquaculture, Salmon farming, Currency markets, Risk management, Hedging, Multi-product hedging, Threshold hedging, GARCH.

Sammendrag

Med begrenset vekstpotensial, kombinert med en risikofylt biologisk produksjon, er behovet for effektiv finansiell risikostyring i norsk lakseoppdrett av økende betydning. I denne oppgaven tar vi for oss to av de viktigste finansielle risikoene lakseoppdrettere må håndtere for å holde seg konkurransedyktige og lønnsomme, pris- og valutarisiko. Derfor undersøker vi et felles laksepris- og valutahedgingproblem for en norsk lakseoppdretter som eksporterer til EU. Først undersøker vi ytelsen til dynamiske strategier modellert med GARCH-modeller, og sammenligner resultatene med den tradisjonelle *naïve hedge*. Vi finner at dynamiske strategier er fordelaktige for lakseoppdrettere, spesielt når det gjelder avkastning. For det andre, undersøker vi hvordan et *multi-product hedging* - rammeverk, som tar potensielle sammenhenger mellom laksepris og valuta i betraktning, presterer sammenlignet med å hedge produktene separat. Vi finner at å ta hensyn til disse sammenhengene gir gode resultater. For det tredje, introduserer vi en ny *threshold* strategi som benytter seg av volatilitetsklynger i laskeprisen. Resultatene våre viser at denne strategien presterer bedre enn de mer tradisjonelle strategiene hva gjelder avkastning, mens den presterer likt når det gjelder variansreduksjon.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The industry of salmon farming emerged in the 1970s in Norway as a government-supported activity to support depressed coastline economies suffering from declining wild fisheries (Liu et al., 2011). Since then, salmon farming has experienced phenomenal growth. With over 478% growth between 1995 and 2019, aquaculture has been the worlds fastest growing food processing industry according to MOWI (2020). Today, Norway is the world's largest producer with over 1.2 million tonnes harvested in 2019, making salmon an important Norwegian export commodity. At the same time, the industry has consolidated from a great number of local family businesses to fewer and larger producers exporting internationally. Today, demand for salmon is steadily increasing, but further industry growth is limited by various factors. Among these, the lack of geographical sites suited for farming is one of the most important. In order to farm salmon, the production site has to satisfy specific requirements such as appropriate sea temperatures, sheltered and protected coastlines, and several biological conditions (MOWI, 2020). Due to a limited number of such sites available worldwide, production mainly takes place in Norway, Chile, Scotland, Canada, and the Faroe Islands (Asche et al., 2013). The lack of suited sea areas is considered to be a major obstacle for further expansion (Hersoug et al., 2021). Furthermore, sea-based farming may harm the surrounding environment due to incidents such as lice, escapes, and diseases. Environmental concerns have therefore resulted in stricter regulations, which limit further growth (Bjørndal and Tusvik, 2019). Limited growth potential and environmental concerns, combined with increasing demand, have led to the development of land-based farming technologies. Such sites are currently in development or already operating in several key markets, adding pressure on sea-based producers (Bjørndal and Tusvik, 2019).

With the described industry development and the inherent biological nature of the production, salmon producers face numerous challenges and risks that affect growth and profitability. This highlights the importance of efficient financial risk management. In this thesis, we address two of the most important risks salmon producers must manage, salmon price and currency exchange rate fluctuations. In particular, we investigate the potential of advanced hedging strategies for the reduction of the producers' exposure to these risks. We analyze the performance of proposed strategies in the context of both individual and joint hedging of these risks. This is done through a stylized case corresponding to a representative Norwegian producer.

The first and most important financial risk we investigate is the salmon price. According to Asche et al. (2018) the salmon price is volatile, especially in comparison to other commodities, and the volatility has more than doubled over the last ten years. Moreover, the salmon price shows volatility clustering effects, meaning extreme values are likely to be followed by more extreme values (Oglend and Sikveland, 2008). Several factors cause the observed volatility. While the three-year production cycle is quite long, the salmon market is mainly a fresh-fish market. Thus, production and consumption must happen within the same period. This means short-term output levels are hard to adjust while demand is affected by factors such as season, quality, and disease outbreaks (MOWI, 2020). Asche et al. (2019) point to this inelasticity as an important cause of the salmon price volatility. In addition, Bergfjord (2007) points to political and regulatory shocks as important determinants of price volatility. High price volatility is troublesome as the salmon price is one of the most critical determinants of the producers' profits and cash flows. Both are heavily affected by fluctuating salmon prices, hence, price risk management is of great importance.

The second financial risk we focus on is the currency exchange rate risk.[1] With large portions of the harvest sold internationally, Norwegian producers have great exposure to changes in exchange rates. Floating exchange rates cause what is known as the importer-exporter dilemma for firms operating in international environments (Berk and DeMarzo, 2014). Both buyers and sellers are affected by exchange rate movements, making the dilemma a general problem for import-export trade, where one of the parties must face the floating rate. In addition to heavily influencing international firms' profits, fluctuating exchange rates also affect the expected future cash flows and thus the value of international firms (Allayannis and Ofek, 2001). All major Norwegian salmon producers state in their annual reports[2] that they have exposure towards a number of currencies and emphasize that fluctuating exchange rates represent a direct financial risk.

In order to manage the price and currency risk, the producers may engage in hedging using futures derivatives. This way, revenues can be secured and risk management costs reduced according to Asche et al. (2016). Commodity futures contracts are traded at regulated exchanges. For salmon derivatives, this exchange is Fish Pool ASA.[3] With regards to currency hedging, most financial institutions offer trading of derivatives. All Norwegian producers listed on the OSLO Seafood Index[4] state in their annual reports, that they engage in hedging using both salmon futures and currency forward derivatives. However, Bloznelis (2016) states that less than 10% of the Norwegian production is hedged through Fish Pool. The low trading volumes indicate that Norwegian producers do not apply advanced hedging strategies as part of their risk management practices. This motivates and forms the basis of our research question:

Can Norwegian Atlantic salmon producers improve current risk management practices by utilizing more advanced strategies for hedging both price and currency risk?

The novelty of this thesis is twofold. First, we introduce a novel application of multi-product hedging where we model the joint risk of price and exchange rate movements using a state-of-the-art DCC-GARCH model. Despite both risks being of substantial importance, no prior studies have looked into the joint hedging of salmon price and exchange rates. According to Haigh and Holt (2002), accounting for all sources of risk is vital when assessing the hedging potential of a particular derivative contract. This

[1] Currency and exchange rate will be used interchangeably throughout the thesis.

[2] MOWI ASA, 2021, SALMAR ASA, 2021, Lerøy Seafood Group ASA, 2021, Grieg Seafood ASA, 2021, Royal Norwegian Salmon ASA, 2022.

[3] Established in 2006 and licensed by the Norwegian Ministry of Finance, Fish Pool ASA aims to create predictability in risk-exposed seafood markets by offering fish and seafood derivatives (Fish Pool, 2020).

[4] https://live.euronext.com/en/product/indices/NO0010760663-XOSL/market-information (Accessed: 2021-03-21)

is because the hedging effectiveness of a new contract may be reduced if other risks are not accounted for, and especially if the prices are correlated. Our first contribution is thus expanding the existing literature on aquaculture hedging, as we are the first to study hedging of both price and currency risk simultaneously in a time-varying setting. Second, we develop a novel threshold hedging strategy that utilizes volatility clustering in the salmon price returns. We compare this self-developed threshold hedging strategy to both static and dynamic hedging strategies in the joint price and currency risk framework. The introduction of this novel strategy complements the existing hedging literature and provides new insights about aquaculture hedging. Lastly, by investigating the performance of different hedging strategies, we provide practical tools for better risk management applicable for salmon producers. However, the hedging strategies and the general methodology are applicable across commodities, and is, therefore, useful for other industries as well.

Our main findings can be summarized as follows. First, we find that more advanced hedging strategies increase returns while performing similarly in terms of risk reduction when compared to the traditional naïve hedge. Second, we find that utilizing dependencies between the salmon and foreign exchange markets through a state-of-the-art multi-product hedging framework is beneficial for salmon producers. Compared to hedging products independently, the multi-product hedge performs similarly in terms of risk reduction but yields higher returns. Third, we find that our self-developed threshold strategy outperforms the other, more traditional strategies in terms of risk-return trade-off. It performs similarly in terms of risk reduction, but yields higher returns. Fourth, we find that it is important to consider the costs of hedging, as it indicates the viability of different hedging strategies for salmon producers. The more advanced hedging strategies are subject to lower transaction costs and, thus, yield superior mean returns compared to the naïve hedge. Fifth, we find that the hedging horizon greatly affects the hedging results. For longer horizons risk reduction is substantially reduced when hedging, but this comes at the cost of decreased returns. Lastly, we find that the development of the salmon prices heavily influences the hedging performance. All considered strategies, except for the naïve hedge, improve both risk reduction and return when prices depreciate. The same is not observed when prices appreciate, indicating that hedging becomes more important during periods with depreciating prices.

The remainder of the paper is organized as follows. First, we review the existing literature in Chapter 2. Then we present the applied methodology in Chapter 3. A description of the data on which we apply the methodology is presented in Chapter 4. Our results are presented and discussed in Chapter 5, before Chapter 6 concludes the paper.

Chapter 2

Literature review

In this chapter, we present and review relevant literature. First, we present literature related to salmon hedging, then currency hedging, and finally, we look at the existing literature within multi-product hedging.

The volatility of the salmon price has been the focus of extensive research. As one of the first studies, Oglend and Sikveland (2008) observe substantial volatility in the salmon price, as well as volatility clustering effects. Later, Oglend (2013) demonstrates empirically that the price of Atlantic farmed salmon from Norway has been increasing since the early 2000s. The steady increase is also confirmed by Bloznelis (2016). Another study by Asche et al. (2018) finds that the price volatility has more than doubled over the last ten years, thus giving more substance to the claim that financial risk management is increasingly important in the aquaculture industry. As a result, there have been conducted numerous studies on how to manage the salmon price risk by hedging with derivatives. Asche et al. (2016) examine the hedging efficiency of Atlantic salmon futures from Fish Pool. They use different methods to obtain optimal hedge ratios (OHR), and find that a bivariate GARCH model performs the best. However, the dynamic models are slightly outperformed by the traditional naïve hedge of hedging one-to-one with one futures contract to every spot contract. They conclude that the use of salmon derivatives may reduce salmon price risk by approximately 30-40%. Similarly, Bloznelis (2018) analyses hedging of the salmon spot price with the use of futures contracts. The study also models the volatility of the spot price using a GARCH framework, and obtains satisfactory hedging results. It concludes that hedging salmon price with futures contracts is a moderately effective way of managing the price uncertainty.

As observed from the mentioned studies, it is common to use GARCH to model the conditional volatility of the salmon price. This is supported by Oglend (2013) who finds strong evidence of heteroscedasticity in the salmon price. Thus, GARCH models are appropriate to use in the context of modeling salmon price uncertainty. According to Brooks et al. (2002), the general consensus is that the use of GARCH models yields superior performance, evidenced by lower portfolio volatility, than either time-invariant or rolling ordinary least squares (OLS). By applying state-of-the art GARCH models we relate to best practice of modern aquaculture literature.

Oglend and Sikveland (2008) use GARCH models to examine volatility clustering in the salmon price. They find the presence and persistence of clustering effects in the salmon price, meaning large price changes are usually followed by more large changes, and small changes are followed by small changes.

This implies that volatility clustering offers predictive information on price fluctuations in the market. Haarstad et al. (2021) observe volatility clustering effects in the salmon spot, but not in the forward price. This asymmetry may provide undiscovered hedging potential. Therefore, in this thesis we extend aquaculture hedging literature by proposing a novel hedging approach that exploit volatility clustering.

Another stream of literature relevant to our thesis deals with currency hedging. The foreign exchange (FX) markets are the largest and most liquid of all asset markets (Chang et al., 2013), and has been the topic of numerous studies. Ødegaard and Børsum (2005), conclude that Norwegian firms are exposed to exchange rate fluctuation, with the most obvious source of impact being import and export prices. Furthermore, their study finds that 91% of Norwegian companies with currency exposure engage in currency hedging. Chang et al. (2013) examine the hedging efficiency obtained by using currency futures. Their results indicate that there are significant GARCH effects in their currency futures, and that a GARCH(1,1) model efficiently explain the uncertainty in the series. Their study concludes that hedging using futures derivatives effectively reduces risk for every currency and maturity considered. Ku et al. (2007) examine different models to decide the OHRs in different currency futures markets. Their study compares the traditional OLS model to the more advanced dynamic conditional correlation (DCC) GARCH model. They find that while both reduces risk effectively, the DCC-GARCH model yields the best hedging performance. Similarly, Chakraborty and Barkoulas (1999) investigate the hedging performance of dynamic strategies using futures contracts for the five biggest currencies. They use a bivariate GARCH model to estimate the joint distribution of spot and futures currency returns, and time-dependent OHRs. While they find that the dynamic hedging model is empirically appropriate, in the case of four out of the five currencies, they do not find significant gains in hedging efficiency compared to the naïve hedge. However, Kroner and Sultan (1993) show that the use of hedge ratios modeled with GARCH yield better hedging efficiency than traditional hedge ratios in currency markets. In this thesis, we follow the state-of-the-art approach and apply the GARCH framework to model the currency risk.

The studies mentioned so far focus on hedging price and currency risk separately. To understand how these risks can be hedged jointly, we investigate literature that considers hedging of commodity price and currency risk simultaneously. As Benninga and Eldor (1985) state, the exporter's hedging problem differs from those generally considered in the literature, as the optimal hedge in one of the markets depend on the size of the hedge in the other market. As all international producers are affected by the commodity price and currency risk, they derive optimal hedging rules for exporting firms. Interestingly, their study finds that the size of the commodity hedge is independent of the properties of the FX market, but the optimal currency hedge depends on the properties of the commodity market. Another paper by Haigh and Holt (2002) examines linkages between freight, commodity and exchange rates. Their results suggest that exploiting co-dependencies between the different products yield improved risk reduction for traders. Yun and Kim (2010) analyzes the hedging effectiveness of different hedging strategies and periods for Korean oil traders, where both crude oil price and exchange rate fluctuations are considered. Their study finds that considering the inter-correlation between the oil price and exchange rate movements improves the hedging effectiveness. In addition, they find that the hedging effectiveness tends to improve as the hedge period increases. Another study by Husodo and Vidiapratama (2011), examines variance reduction by conducting cross-hedging for a US grain trader, with both commodity price and currency exposure. They find that multivariate GARCH models better describe the joint dynamic behaviour of commodity prices and currency rates, compared to conventional models. Their results show that the

GARCH models outperform the OLS model in finding the optimal hedging strategy in most markets. Nevertheless, they emphasize the importance of considering the commodity price and currency rate jointly, as considering them in isolation ignores potential co-dependencies.

Despite the fact that there are a substantial body of research that investigate joint hedging of commodity price and currency risk, there are no prior studies to our knowledge that investigate this with application to the salmon market. Among the first contributions that explicitly consider multi-product hedging in the context of salmon farming was a study by Haarstad et al. (2021). They consider different multi-commodity hedging strategies for the joint input and output price risk for salmon producers. Their results indicate great potential for hedging efficiency in the context of multi-product hedging within aquaculture. We contribute to this literature by filling the gap by introducing exchange rate to a multi-product hedge.

To summarize, we find numerous studies that investigate hedging of salmon price and exchange rate risk separately, but none studies that hedge these risks jointly. We extend the aquaculture literature by investigating the potential benefits of joint salmon price and currency hedging strategies, and by introducing a novel hedging strategy that utilizes volatility clustering effects. Overall, we place ourselves among a limited set of researchers that conduct this multi-product hedging analysis.

Chapter 3

Methodology

In this chapter, we introduce the methodology we apply in order to investigate the potential of advanced hedging strategies for Norwegian salmon producers. First, we provide a detailed description of the assumptions behind our case study in Section 3.1. Following this, we introduce a set of hedging strategies in Section 3.2. We then describe how we model the uncertainty in the time series using the GARCH framework in Section 3.3, and finally how we evaluate the hedging performance in Section 3.4.

3.1 Case description

Our goal is to investigate if producers of Atlantic farmed salmon can improve their financial risk management practices through advanced hedging strategies. In order to do this, we consider a hypothetical Norwegian producer of Atlantic farmed salmon that corresponds to a large Norwegian producer. This well-established salmon producer aims to reduce its exposure to fluctuations in the salmon price and the currency exchange rate. With this in mind we make the following assumptions.

First, we assume the producer is mainly located in Norway and reports its earnings in Norwegian kroner (NOK). Therefore, we assume that the company harvests quantities similar to the average of the largest Norwegian salmon producers on the Oslo Seafood Index.[1] Using the data from the companies' 2020 annual reports, we calculate a harvest of approximately 180 000 tonnes a year, or a weekly average of about 3500 tonnes. This harvest volume is similar to the harvest volumes of companies such as Salmar and Lerøy Seafood Group.

Second, we assume that the producer harvests and sells the salmon continuously throughout the year, with an average weekly volume of 3500 tonnes. All harvested salmon is sold every week, i.e. no carryover inventory to the next week. In addition, we assume that all salmon sold within the same week achieve the same price.

The company aims to reduce the exposure to fluctuations in the salmon spot price. This is achieved through the use of futures contracts from Fish Pool. We use a one month contract length as the front month forward price is most correlated with the spot price.[2] In addition, using a shorter contract

[1] https://live.euronext.com/en/product/indices/NO0010760663-XOSL/market-information (Accessed: 2021-03-09)

[2] Longer contracts, 3, 6, and 12 months, do not exhibit sufficient correlation with the spot price to be used as efficient hedging tools.

length is reasonable because sales volumes are harder to predict the longer the horizons are. This is in accordance with Asche et al. (2016) and Haarstad et al. (2021). However, studies have found that the Fish Pool futures contracts suffer from liquidity issues due to low trading volumes and participation (Oglend, 2013; Bloznelis, 2018; Oglend and Straume, 2019). We will not consider this problem in our model, and we assume that all futures contracts can be initialised and liquidated simultaneously every week, in order to focus on comparative performance of different hedging strategies.

Next, we assume the producer exports a substantial part of the harvested volume. This implies that the revenue is realized in several different currencies. We have chosen to solely focus on export to the European Union (EU) as this is the most important market for Norwegian producers. In accordance with industry numbers described in MOWI (2020), we assume that 50% of the total harvested volume is sold in euro (EUR). The volume sold in EUR is therefore 1750 tonnes weekly. Thus, the producer has exposure to both salmon price and exchange rate fluctuations. In order to hedge the exchange rate fluctuations, we assume that the producer can buy exchange rate derivatives from a Norwegian financial institution.

In the baseline scenario, we use a four-week hedging horizon. This is common practice in comparable studies that look into hedging in the salmon farming industry, such as Haarstad et al. (2021) and Asche et al. (2016). However, we verify the sensitivity of our results with respect to different horizons.

Lastly, we explicitly account for transaction costs in the trading of salmon futures contracts by including a fixed fee for every transaction. The fee is set to 0.15 NOK/kg, in accordance with standard Fish Pool contracts.[3] This includes both clearing and trading. For the currency contracts, we assume no extra transaction cost as this is usually already incorporated in the forward rates.[4]

Based on these assumptions, we define the following two-period hedged portfolio return for the salmon sold in EUR and converted to NOK:

$$\begin{aligned}\pi_t(h) \;=\; & Q^{SA}(S_t^{SA}S_t^C - S_{t-i}^{SA}S_{t-i}^C) \;-\; h^{SA}Q^{SA}(F_t^{SA}F_t^C - F_{t-i}^{SA}F_{t-i}^C) \\ & -\; h^C Q^{SA}(F_t^C - F_{t-i}^C)F_{t-i}^{SA} \quad .\end{aligned} \tag{3.1}$$

Here, superscript SA and C denotes salmon and currency, respectively. Q^{SA} is the weekly quantity of salmon sold, S_t and F_t denotes the spot and forward prices when the hedges are liquidated, and similarly S_{t-i} and F_{t-i} when the hedges are initialised. Finally, h denotes the hedge ratios. This portfolio definition is consistent with similar studies on joint commodity price and currency hedging, such as Yun and Kim (2010) and Haigh and Holt (2002).

3.2 Hedging strategies

3.2.1 Static and dynamic hedging strategies

A producer may hedge by taking opposite positions in the spot and futures market for salmon and/or currency. This way, an adverse fluctuation in either market can be offset by a favorable counter-movement in the other. To decide on the composition of the hedged portfolio, a hedge ratio denoted h, is defined. The hedge ratio is the number of units of futures contracts purchased relative to the exposure

[3] https://fishpool.eu/trading/fee-list/ (Accessed: 2021-04-07)

[4] This is common practice for financial institutions, as exemplified by DNB. (https://www.dnb.no/en/business/markets/foreign-exchange/hedging/forwards.html (Accessed: 2021-03-10))

in the spot market (Brooks et al., 2002). In other words, the hedge ratio is the size of the position in the futures market. The most common measure of risk in this context is the variance of the hedged portfolio. The objective is therefore to find the hedge ratio that minimizes the variance of the portfolio returns. This so-called minimum variance (MV) hedge ratio is simple to understand and estimate (Lee et al., 2003). Different strategies can be applied for choosing the hedge ratio. In the following, we look at static and dynamic hedging.

A strategy where the hedge ratio is kept constant over the hedging horizon is known as static hedging. Choosing a hedge ratio for a static hedge can be done using two approaches. One simple approach is to hedge the risk by taking one unit of a short position of a futures contract for each unit of a long position in the spot. This is known as the *naïve hedge*, with $h = 1$ (Wang et al., 2015). An implicit assumption of the naïve approach is that the spot and futures prices move closely together, and that a perfect hedge can only be achieved if proportionate price changes in one market exactly matches those in the other market (Butterworth and Holmes, 2001). This approach is simple to implement and control, and is often reasonable to use if there is a lack of information.

The second static approach is to calculate the optimal hedge ratio (OHR), h^*, which minimizes the variance of the hedged portfolio returns. This approach assumes that the joint distribution of spot and futures returns is time-invariant (Chang et al., 2013). Unlike the naïve hedge, the OHR approach does not require perfect correlation between the spot and futures markets in order for the hedge to be optimal, as it accounts for imperfect correlations. However, as it is a static hedging strategy, the OHR is estimated under the assumptions of constant volatility and correlation. The only static hedging strategy we consider is the naïve hedge, which we use as a benchmark, as our main focus is on the performance of the more advanced strategies.

A strategy where the hedge ratio can change over the hedging horizon is known as dynamic hedging. The variance and covariance, and therefore the correlations of asset returns, are time-varying (e.g., Bollerslev et al., 1988; Engle, 2002, among many others). This implies that the OHRs determined by these variances and covariances also are time-variant (Wang et al., 2015). With the dynamic hedging strategy there is no need for us to assume constant volatility and correlation, which is often unrealistic in financial time series. The objective is to find the optimal time-varying hedge ratio at time t. To calculate the OHR we let $r_{s,t}$ and $r_{f,t}$ denote the returns of spot and futures prices at time t, respectively. Let h_t be the hedge ratio at time t. The return of the hedged portfolio at time $t+1$, denoted $r_{p,t+1}$, is then given by the following equation:

$$r_{p,t+1} = r_{s,t+1} - h_t r_{f,t+1} \quad . \tag{3.2}$$

The variance of the hedged portfolio is given by:

$$Var(r_{p,t+1}) = Var(r_{s,t+1}) + {h_t}^2 Var(r_{f,t+1}) - 2h_t Cov(r_{s,t+1}, r_{f,t+1}) \quad . \tag{3.3}$$

By minimizing this equation, we derive the following MV dynamic hedge ratio:

$${h_t}^* = \frac{Cov(r_{s,t+1}, r_{f,t+1})}{Var(r_{f,t+1})} \quad . \tag{3.4}$$

Here $Cov(r_{s,t+1}, r_{f,t+1})$ is the conditional covariance between the returns of the spot and futures prices, and $Var(r_{f,t+1})$ is the conditional variance of the futures returns (Wang et al., 2015; Brooks, 2014). In order to find the OHR using this strategy we have to estimate these two measures. Estimation procedures are discussed in Section 3.3.

3.2.2 Single- and multi-product hedging

Different approaches may be applied when considering multiple products in a hedging context. In this study, we consider both a single-product hedge and a multi-product hedge. We start by defining a single-product hedge, henceforth referred to as *single hedge*, where products are considered separately (in our setting, salmon and currency) without exploiting potential dependencies between the two. The OHR for a single hedge is calculated using the framework outlined in Section 3.2.1. The return of the hedged company portfolio is then a combination of the MV portfolio of each product considered and hedged independently using Equation 3.4. This is the current practice in the salmon farming industry today, where hedging of different products is done independently of each other.[5] An advantage with this approach is that it is simple. However, one might miss out on the potential upside resulting from correlation between the salmon and currency markets.

In order to exploit dependencies between the different products, we define multi-product hedging, henceforth referred to as *multi hedge* The return on the company portfolio is considered in a multi-product setting, which in our case means that the salmon and currency are considered in unison, and not independently as with the single hedge. Given that there exist dependencies between the products, this approach implies that adverse fluctuations in the price of one product can be offset by favorable movements in the price of the other. For instance, unfavorable movements in the salmon spot price could be offset by movements in the exchange rates, rather than just by the salmon futures price. While the OHR calculated using Equation 3.4 holds for the single hedge approach, it might not be optimal when considering a multi-product problem. Therefore, we use the framework formalised by Fackler and McNew (1993) for finding optimal hedge ratios in situations with multiple spot and futures series. The vector of optimal time-varying hedge ratios, denoted $h^*_{M,t}$, is given by the following equation:

$$h^*_{M,t} = [diag(Q)]^{-1} \sum\nolimits_{FF}^{-1}(t) \sum\nolimits_{FS}(t) Q \quad . \tag{3.5}$$

In this equation, Q is an m - vector of the quantities of spot products, with positive signs denoting long positions in the product and negative denoting short positions. m is the number of products considered. As we consider one commodity and one currency, $m = 2$ in our case. $diag(Q)$ is a diagonal matrix with the vector Q on its diagonal, $\sum_{FF}(t)$ is the $(m \times m)$ time-varying variance-covariance matrix of futures prices and $\sum_{FS}(t)$ is an $(m \times m)$ matrix representing the time-varying covariances between the spot and future prices. The original framework formalised by Fackler and McNew (1993) has been extended from the static case to the time-varying in Equation 3.5.

As the single hedge and multi hedge apply time-varying OHRs, we refer to these strategies as *dynamic strategies*.

[5] This insight was revealed in a phone interview with an industry representative.

3.2.3 Threshold hedging strategy

In addition to the strategies typically considered in the literature, we introduce a novel strategy, which we call *threshold strategy*. This strategy is based on differences in the volatility of salmon spot and forward returns. It exploits the possible predictive power provided by the observed volatility clustering in the salmon spot price (Oglend and Sikveland, 2008), and the fact that it is not observed in the forward price (Haarstad et al., 2021). The idea of the strategy is, therefore, to initiate a hedge if a spike in the returns of the salmon spot price is observed. The condition used to identify such spikes is if the returns over the past month is higher than a threshold, which is defined as the long-run average of the return series. The condition is specified as follows:

$$\left| \frac{1}{4} \sum_{i=1}^{4} r_{t-i}^{SA} \right| > \bar{r}^{SA} . \tag{3.6}$$

Here, r_{t-i}^{SA} is the salmon spot return at time $t - i$, and $\bar{r}^{SA}$ is the long-run average of the salmon spot return series. If this condition is satisfied, a naïve hedge is initiated. Hence, we refer to it as the *naïve threshold hedge*.

The threshold strategy can be viewed as a form of dynamic hedging strategy as the hedge ratio can change over the hedging horizon. However, it differs from traditional strategies because it incorporates a new decision rule that identifies if it is worth to initiate a hedge in the first place. Hence, by utilizing this strategy a decision maker is provided with a hedge ratio and guidance on when to hedge. In that way, this strategy is more advanced than a simple naïve hedge. However, as a simple decision rule is provided the threshold strategy should be realistic to implement for salmon producers. As the strategy is based on the expected predictive power of volatility clustering in the salmon spot returns, a potential advantage is that it manages to hedge some of the most volatile returns, which could lead to lower portfolio variance. Another potential advantage is lower transaction costs, as the strategy is inclined to take on fewer futures positions.

In addition to using the naïve threshold hedge, we also introduce a more advanced threshold strategy that utilizes the dynamic OHR from the single hedge described earlier. We do this for two main reasons. First, we want to investigate if we can improve the performance of the naïve threshold hedge by utilizing a dynamic hedge ratio. Second, we do this to investigate if we can improve the performance of a traditional dynamic hedge by introducing a threshold decision rule. Henceforth, we refer to this strategy as the *dynamic threshold hedge*.

As the FX market is more efficient, we do not observe the same volatility characteristics as for the salmon market. Therefore, we do not hedge the currency exposure utilizing the threshold strategies.

3.3 GARCH modeling

In this section, we present further details about the methodology behind the modeling of the uncertainty in our financial time series. We apply the generalized autoregressive conditional heteroscedasticity (GARCH) methodology of Bollerslev (1986), which allows for heteroscedasticity in the time series, i.e. time-varying volatility. In GARCH models, the current period's conditional variance is a function of its own previous lags. Below we present more details about the univariate GARCH model GARCH(1,1) and the multivariate GARCH model DCC–GARCH.

First, we define the conditional values of the mean and variance of a time series. Consider a time series of asset prices where the continuously compounded return, defined as the log-return, at time t is denoted y_t. We denote the unconditional mean and variance of the time series μ and σ^2, respectively. The conditional values, μ_t and ${\sigma_t}^2$, given the set of information ψ_{t-1} at time $t-1$, are then given by:

$$\mu_t = E[y_t|\psi_{t-1}] \ , \tag{3.7}$$

$${\sigma_t}^2 = E[(y_t - \mu_t)^2|\psi_{t-1}] \ . \tag{3.8}$$

We can now use these definitions to further define the mentioned GARCH models.

3.3.1 Univariate GARCH

Following the framework introduced by Bollerslev (1986), however with different notation, we let z_t denote a real valued stochastic process and, again, ψ_{t-1} the set of information available at time $t-1$. The univariate GARCH(p,q) process is then given by the following:

$$z_t|\psi_{t-1} \sim N(0, {\sigma_t}^2) \ , \tag{3.9}$$

$$z_t = \sigma_t \epsilon_t \ , \tag{3.10}$$

$${\sigma_t}^2 = \alpha_0 + \sum_{i=1}^{q} \alpha_i z_{t-i}^2 + \sum_{i=1}^{p} \beta_i \sigma_{t-i}^2 \ , \tag{3.11}$$

with the following constraints:

$$\begin{aligned} p &\geq 0 \ , \\ q &> 0 \ , \\ \alpha_0 &> 0 \ , \\ \alpha_i \geq 0 & \qquad i = 1, \ldots, q \ , \\ \beta_i \geq 0 & \qquad i = 1, \ldots, p \ . \end{aligned}$$

Here, ϵ_t is the standardised residual at time t, which is assumed i.i.d. Usually the standard normal distribution is applied, i.e. $\epsilon_t \sim N(0,1)$, like in the original model by Bollerslev (1986). However we also apply the generalized error distribution (GED) and the skewed Student's t distribution. We see from Equation 3.11 that the conditional variance ${\sigma_t}^2$ is dependent on its own past values σ_{t-i}^2, and the value of p decides the number of lags of the variance to include. The unconditional variance of z_t is given by:

$$var(z_t) = \frac{\alpha_0}{1 - (\sum_{i=1}^{q} \alpha_i + \sum_{i=1}^{p} \beta_i)} \ . \tag{3.12}$$

In order to ensure stationarity in the process we require that $\sum_{i=1}^{q} \alpha_i + \sum_{i=1}^{p} \beta_i < 1$ (Palm, 1996; Brooks, 2014).

GARCH(1, 1) model

With the above framework it is possible to specify a number of different models through the index (p, q). However, the simple GARCH$(1, 1)$ model has been found to be sufficient in most financial time series and is widely used in the literature (Hansen and Lunde, 2005; Brooks, 2014; Palm, 1996). Therefore, we use the GARCH$(1, 1)$ model on the data in our study. Based on the model described above, the $(1, 1)$ specification gives us the following conditional variance:

$$\sigma_t{}^2 = \alpha_0 + \alpha_1 z_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \ . \tag{3.13}$$

In this model $z_t{}^2$ and $\sigma_t{}^2$ are lagged from one period before. Stationarity is ensured when $\alpha_1 + \beta_1 < 1$.

3.3.2 Multivariate GARCH

The framework outlined in the previous section models the volatility for a single time series. However, as we are interested in the dependencies between multiple time series, we need to employ a multivariate GARCH model in order to capture the time-varying co-movements. In the following, we present the standard multivariate GARCH framework. Next, we present the DCC-GARCH model, which is the multivariate model we apply on our data. We follow the same setup as Silvennoinen and Teräsvirta (2009). Consider a stochastic process vector $\{\mathbf{z_t}\}$ with dimension $(N \times 1)$. As earlier, we let ψ_{t-1} denote the set of available information up to and including time $t - 1$. $\{\mathbf{z_t}\}$ is assumed to be conditionally heteroscedastic, and given by:

$$\mathbf{z_t} = \mathbf{H}_t{}^{1/2} \epsilon_t \ , \tag{3.14}$$

$$\mathbf{z_t} | \psi_{t-1} \sim N(0, \mathbf{H}_t) \ , \tag{3.15}$$

where $\mathbf{H}_t$ is an $(N \times N)$ - matrix of the conditional covariances of $\mathbf{z_t}$ and ϵ_t is an $(N \times 1)$ i.i.d. error process vector with the properties $E(\epsilon_t) = 0$ and $Var(\epsilon_t) = I$. In our case, $\mathbf{z_t}$ is a vector of the log-returns of the N time series we look at.

This defines the standard multivariate GARCH framework. What remains, is to specify the covariance matrix $\mathbf{H}_t$. There are a number of different specifications in the literature, see for instance Wang et al. (2015) or Bauwens et al. (2006). The two most widely used models for modeling conditional covariances and correlations are the BEKK model proposed by Engle and Kroner (1995), and the dynamic conditional correlation (DCC) model proposed separately by Engle (2002) and Tse and Tsui (2002). The two models are similar in many ways. However, the BEKK model suffers from the curse of dimensionality, meaning that the number of parameters in the model increase at an order higher than the number of assets (Caporin and McAleer, 2009). Therefore, we use the DCC model in this study.

DCC-GARCH model

The DCC model is a generalization of the constant conditional correlation (CCC) model proposed by Bollerslev (1990). The only difference between these models is that the DCC model relaxes the often unrealistic assumption of time-invariance in the conditional correlations over time. For these models the conditional variances of the spot and futures series follow the univariate GARCH(1,1) model, which simplifies the estimation of the conditional covariances. The variance-covariance matrix, $\mathbf{H}_t$, for the DCC model is defined as follows:

$$\mathbf{H_t} = \mathbf{D_t R_t D_t} \ , \tag{3.16}$$

where $\mathbf{D}_t$ is a diagonal matrix of the time-varying standard deviations from the estimation of the univariate GARCH(1,1) using Equation 3.11, and $\mathbf{R}_t$ is the conditional correlation matrix (Creti et al., 2013). As indicated with the subscript t, both $\mathbf{D}_t$ and $\mathbf{R}_t$ are time-variant, which means the variance-covariance matrix changes with each time step. We apply a multivariate normal distribution when modeling our time series with DCC.

With the DCC model we estimate the variance-covariance matrices needed for calculating the optimal hedge ratios in the single hedge and multi hedge frameworks mentioned in Section 3.2.2.

3.4 Hedge performance evaluation

To evaluate and compare the performance of the hedging strategies, we need a measure of hedging efficiency. A common and widely used measure of risk reduction is the hedge effectiveness (HE) measure proposed by Ederington (1979). HE is measured as the reduction in the variance of the hedged portfolio compared to the variance of the unhedged portfolio, and is given by the following equation:

$$HE = \frac{var_{unhedged} - var_{hedged}}{var_{unhedged}} \ . \tag{3.17}$$

Another risk measure is the expected shortfall (ES), proposed by Acerbi et al. (2001). ES is a measure of tail risk in a portfolio, which is the risk of major losses occurring due to extreme events. Value-at-risk (VaR) is the most widely used measure of tail risk, but it is not a coherent risk measure as it does not possess the property of sub-additivity. This means the total portfolio VaR can be larger than the sum of the sub-portfolio VaRs. This brings a computational disadvantage since our portfolio, as presented in Equation 3.1, in essence is made up of two sub-portfolios, i.e. salmon and currency. In addition, VaR is indifferent to the severity of the worst case losses (Acerbi and Tasche, 2002). Due to these shortcomings we use ES to measure the tail risk. ES is measured as the average of the $\alpha = A\%$ worst losses, and given by the following equation:

$$ES^{(\alpha)}(X) = -\frac{1}{\alpha}\Big(\mathbb{E}[X\mathbf{1}_{\{X \leqslant x^{(\alpha)}\}}] - x^{(\alpha)}(\mathbb{P}[X \leqslant x^{(\alpha)}] - \alpha)\Big) \ . \tag{3.18}$$

In addition to looking at the hedging performance from the risk perspective, we also look at it from the perspective of returns. The mean returns of the differently hedged portfolios are calculated, based on historical returns. This provides an indication of how the different hedging strategies would have performed over the period of the historical data used. Lastly, we calculate the transaction costs for the different hedging strategies. This is useful as the different strategies yield different hedge ratios, and therefore different positions. The transaction costs provide useful insights into the costs of hedging.

Chapter 4

Data

In this chapter, we present the data used in this study and examine its characteristics. In addition, we present the estimated GARCH models. Section 4.1 focuses on the price and return series for the salmon prices and currency exchange rates, as well as on the suitability of the GARCH framework for modeling the uncertainty in returns. In Section 4.2, we present the estimated GARCH models.

4.1 Data series

Our data set consists of four time series: (1) salmon spot price [EUR/kg], (2) salmon one month forward [EUR/kg], (3) currency exchange rate spot [NOK/EUR] (4) currency exchange rate one month forward [NOK/EUR]. Salmon spot prices are collected from Bloomberg Market Data Feed and represent the Fish Pool Index (FPI). The FPI is a reference price based on a weekly average price for 3-6 kg superior quality, head-on gutted salmon.[1] The salmon forward prices are downloaded directly from Fish Pool.[2] Our exchange rate data are gathered from Refinitiv Eikon Datastream.[3]

Our data set consists of 624 observations of weekly prices collected in the period 01.01.2008 to 31.12.2019.[4] The starting date is the first date Fish Pool offered references prices for both spot and forward prices, while the end date was chosen to exclude the effects of the Covid-19 pandemic. We eliminate one observation at the end of 2010 from our time series, which we consider an outlier.[5] We divide the data into two sub samples where we use in-sample data for estimation and out-of-sample data to test the model on unknown data. The in-sample data consists of 468 observations from 01.01.2008 to 31.12.2016 for each of the time series. The out-of sample data consists of 156 observations from 01.01.2017 to 31.12.2019.

Figure 4.1 shows the salmon spot and forward prices, and the spot and forward exchange rates. Looking at the historical development of the salmon prices in Figure 4.1a, we observe several interesting features, including frequent fluctuations, seasonality and sudden price drops. With prices frequently fluctuating

[1] https://fishpool.eu/price-information/spot-prices/ (Accessed: 2021-04-03)

[2] https://fishpool.eu/price-information/forward-prices-3/forward-closing-prices-history/# (Accessed: 2021-04-03)

[3] Spot rates are based on median rates from the Refinitiv Market Data System, sampled in a five minute window around 16:00 each day. The one month forward rate is collected by Refinitiv Eikon.

[4] We adjust the daily salmon prices into weekly by using the settlement price on the last business day of the week.

[5] This outlier could be explained by strong seasonal demand and limited trading during the last week of December (Asche et al., 2016).

between 3 EUR/kg at the lowest to almost 9 EUR/kg at the highest, the need for efficient price risk management is apparent. Some of the largest price movements can be attributed to major one-time events, such as the 2009 Chilean disease crisis. This crisis led to substantially lower harvest volumes and is the most probable cause of the subsequent increase in salmon prices (Asche et al., 2009). Figure 4.1b shows the exchange rates, and we observe that they are not as volatile as the salmon price, indicating a more stable and mature market. We also observe that the spot and forward seem to move closely together, indicating an efficient market.

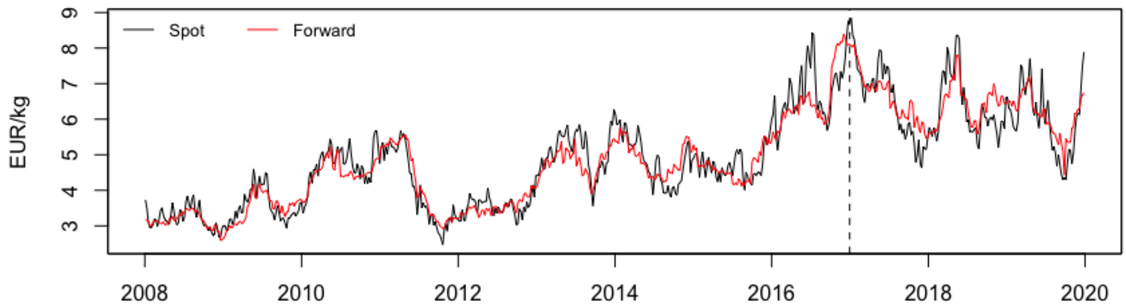

(a) Salmon prices

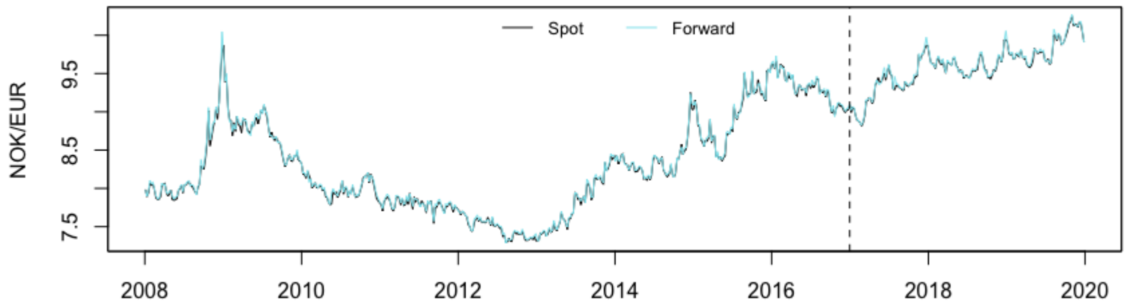

(b) Exchange rates

Figure 4.1: Salmon spot and 1M forward [EUR/kg], and exchange rates spot and 1M forward [NOK/EUR].
Note: Dashed line divides in- and out-of-sample data.

As mentioned in Section 3.3, we use the continuously compounded returns when estimating the volatility. Figure 4.2 illustrates the weekly log-returns for our four time series, and Table 4.1 presents descriptive statistics for the in-sample returns. Descriptive statistics for the out-of-sample returns can be found in Table 4.2.

Table 4.1: In-sample descriptive statistics for weekly spot and forward returns.

	Mean	Median	Min.	Max.	SD	Skewn.	Kurt.
Salmon spot	0.0018	0.0024	-0.1859	0.1570	0.0630	-0.1036	-0.0678
Salmon forward	0.0020	0.0024	-0.1545	0.1294	0.0346	-0.2886	2.3823
Currency spot	0.0003	0.0001	-0.0523	0.0466	0.0119	0.2555	2.8117
Currency forward	0.0003	-0.0002	-0.0516	0.0596	0.0122	0.4070	3.2170

We observe from Figure 4.2a that the salmon spot returns are more volatile than the forward. This is confirmed by the descriptive statistics in Table 4.1 where the standard deviation is higher for the spot series, which is in line with the Samuelson hypothesis (Samuelson, 1965). Furthermore, we see that the mean is positive for both salmon returns series, meaning that salmon returns, on average, appreciate over time. While both the salmon spot and forward return series are negatively skewed, the kurtosis of the forward is high (>2) compared to the spot, which means there is a greater chance of extreme negative values. Thus, fluctuations of the forward series are likely to be negative compared to the expected returns. The characteristics of the salmon returns observed here will have implications for the GARCH modeling, especially for the fitted distributions. The spot series is close to normally distributed with mean, skewness and kurtosis close to zero. However, the forward series does not seem to follow a normal distribution. This is confirmed by the Jarque-Bera test (JB) developed by Jarque and Bera (1980). The JB test results are presented in Table 4.3, and we observe that the null hypothesis of normality is rejected for the forward, but not for the spot.

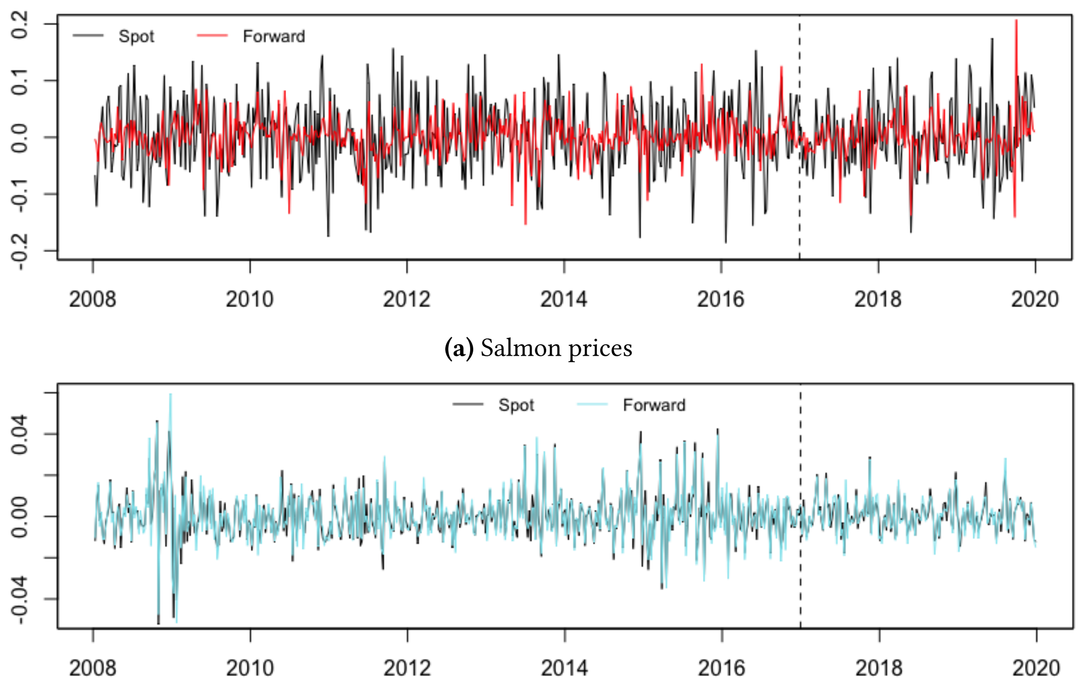

(a) Salmon prices

(b) Exchange rates

Figure 4.2: Weekly log-returns for salmon spot and 1M forward, and the spot and 1M forward exchange rates.
Note: Dashed line divides in- and out-of-sample data.

From Figure 4.2b, we observe that the exchange rate returns seem to follow each other closely, indicating that they are highly correlated. It implies that the forward is an efficient tool for hedging the spot exposure. From the descriptive statistics we see that the mean of both exchange rate return series are close to zero, although slightly positive. This indicates that they marginally appreciate over time. When it comes to the volatility, we observe that they are quite close, again indicating that the spot and forward move together. Next, we see that the kurtoses for both series are quite high, hence, the distributions are fat-tailed. This indicates non-normality, which is also confirmed by the JB test results in Table 4.3. The positive skewness indicates that the most extreme returns are likely to be positive, implying asymmetry in the exchange rate exposure. Findings by Lien (2009) suggest that higher skewness induce higher positions in the futures markets, thus we expect higher hedge ratios than for normally distributed returns.

The out-of-sample descriptive statistics for the salmon returns in Table 4.2 differs substantially from the in-sample descriptive statistics. We notice that the salmon spot and forward returns have negative mean return values. Unlike the in-sample returns, this implies that the salmon prices, on average, depreciate over time. With depreciating prices, it is increasingly important to secure positive returns, and thus to identify the strategies that perform well in terms of return and risk. For the exchange rates the characteristics of the out-of-sample returns are quite similar as for the in-sample returns.

Table 4.2: Out-of-sample descriptive statistics for weekly spot and forward returns.

	Mean	Median	Min.	Max.	SD	Skewn.	Kurt.
Salmon spot	-0.0006	-0.0071	-0.1686	0.1751	0.0633	0.2356	-0.1854
Salmon forward	-0.0012	-0.0029	-0.1408	0.2072	0.0392	0.3947	6.2024
Currency spot	0.0006	0.0000	-0.0178	0.0289	0.0078	0.6102	0.9521
Currency forward	0.0005	0.0001	-0.0190	0.0283	0.0084	0.3905	0.5841

The GARCH framework relies on the assumption of stationarity. Visual inspection of Figure 4.2 suggests stationarity, which is confirmed by the Augmented Dickey-Fuller (ADF) test for unit roots with lag length according to Schwert (1989).[6] The test is applied to the return series with the null hypothesis that they are non-stationary. As seen in Table 4.3, the null is strongly rejected for all series. In addition, a KPSS[7] test with no drift and no trend is conducted to verify the results of the ADF test. The null hypothesis for the KPSS test is that the data is stationary and is not rejected for any of the return series.

Table 4.3: In-sample test statistics for weekly spot and forward returns.

	JB	ADF	LBQ	LM
Salmon spot	0.93	-5.10 ***	5.65 (0.686)	31.2***
Salmon forward	116.92***	-4.90***	1.51 (0.993)	94.1***
Currency spot	158.91***	-5.69***	11.02 (0.201)	72.6***
Currency forward	214.27***	-5.65***	9.53 (0.300)	75.8***

Note: Statistics based on significance levels of $p^{***} < 1\%$, $p^{**} < 5\%$ and $p^{*} < 10\%$

After stationarity is verified, we test for autoregressive conditional heteroscedasticity (ARCH) effects in the return series, to evaluate the appropriateness of using GARCH. We fit the return series to autoregressive (AR) models with order based on the autocorrelation function (ACF), partial ACF (PACF) and the Akaike Information Criterion (AIC). We use the residuals from the optimal AR models to test for autocorrelation by conducting a Ljung-Box q test (LBQ). The results show that the null cannot be rejected for a lag length up to eight, indicating absence of autocorrelation in the returns. Furthermore, we use Engle's Lagrange Multiplier (LM) test to investigate whether or not there are ARCH effects in all of the residual series. The null hypothesis of no ARCH effects are strongly rejected for all series using a lag length of 12.

[6] $lag = 100(\frac{T}{100})^{\frac{1}{4}}$, T = no. of observations.
[7] Kwiatkowski–Phillips–Schmidt–Shin.

The out-of-sample test statistics results are presented in Table A.1 and show the same characteristics as for the in-sample results. We conclude that GARCH is suitable for modeling the returns in- and out-of-sample.

4.2 Estimated models

In this section, we present the estimation results for our GARCH models. We start with the univariate models, which we then use in the estimation of the multivariate models. The estimated GARCH(1,1) models for each time series are presented in Table 4.4. The table shows estimated model parameters with their robust standard errors and p-values. This is followed by the skew and shape parameters for the fitted distributions. Lastly, test statistics and p-values for the weighted ARCH LM test and the Adjusted Pearson Goodness-of-Fit test are presented.[8]

From Table 4.4 we can see that all model parameter estimates are significant at the 5% level, except for the α - parameter for the salmon spot and forward series. In other words, we cannot reject the null hypothesis that $\alpha = 0$. However, if we use a 10% significance level, we see that the salmon spot series is significant while we still cannot reject the null for the forward. This implies that short-term shocks have little impact on the volatility of the salmon forward series, meaning that we do not observe volatility clustering. Visual inspection of Figure 4.2 also suggests absence of volatility clustering in the forward returns, while the spot returns seem to exhibit it to some degree. This justifies the use of the threshold strategy outlined in Section 3.2.3.

The weighted ARCH LM test, based on the theoretical framework by Li and Mak (1994), is a test of adequacy in the fitted GARCH model (Fisher and Gallagher, 2012). From the results displayed in Table 4.4 we see that all models sufficiently capture ARCH effects at the 5% level for all return series and all lags.

The Adjusted Pearson Goodness-of-Fit test proposed by Vlaar and Palm (1993) compares the empirical distribution of the residuals with the theoretical distribution. The results presented in Table 4.4 indicate that all distributions are adequately specified. This is also confirmed by the QQ plots and the plots of the empirical density of standardized residuals presented in Appendix A.2. The most suitable distribution for each return series is chosen through comparison of APG-o-F results and QQ plots for various distributions for all series.

[8] Data plots for the estimated models can be found in Appendix A.2, namely conditional standard deviation (vs |returns|), empirical density of standardized residuals and QQ Plots.

Table 4.4: Estimated GARCH(1,1) models

	Salmon spot			**Salmon forward**			**Currency spot**			**Currency forward**		
Dist.	Normal			GED			Skewed std			Skewed std		
Model par.	Est.	Std. Err.	p-value	Est.	Std. Err.	p-value	Est.	Std. Err.	p-value	Est.	Std. Err.	p-value
α	0.0631	0.0329	0.0554	0.0713	0.0434	0.1004	0.0995	0.0280	0.0004	0.0975	0.0289	0.0007
β	0.6415	0.1265	0.0000	0.7702	0.0803	0.0000	0.8558	0.0426	0.0000	0.8425	0.0458	0.0000
ω	0.0012	-	-	0.0002	-	-	0.0009	-	-	0.0000	-	-
Skew	-	-	-	-	-	-	1.1691	0.0790	0.0000	1.1577	0.0633	0.0000
Shape	-	-	-	1.0498	0.0830	0.0000	6.8938	2.4072	0.0042	5.6963	1.5688	0.0003
ARCH LM	Stat.		p-value	Stat.		p-value	Stat.		p-value	Stat.		p-value
Lag[3]	0.1583		0.6908	0.5273		0.4677	1.549		0.2133	1.869		0.1716
Lag[5]	1.4059		0.6174	1.0843		0.7081	5.092		0.0978	5.156		0.0946
Lag[7]	2.0695		0.7026	1.5094		0.8195	7.564		0.0655	7.123		0.0817
APG-o-F	Stat.		p-value	Stat.		p-value	Stat.		p-value	Stat.		p-value
Group 20	16.30		0.6373	27.00		0.1046	16.04		0.6546	19.55		0.4220
Group 30	29.64		0.4322	37.73		0.1284	21.03		0.8581	35.55		0.1871
Group 40	37.41		0.5424	43.92		0.2709	43.24		0.2952	51.29		0.0900
Group 50	57.30		0.1943	66.08		0.0522	38.67		0.8550	50.67		0.4076

Notes: **1.** Model parameter estimates highlighted in grey are not statistically significant at the 5% level. **2.** The ARCH LM test is conducted with a null hypothesis of no presence of ARCH effects. **3.** The Adjusted Pearson Goodness-of-Fit test compares the theoretical distribution with the empirical distribution, with a null that the two are identical.

Figure 4.3 presents the modeled conditional volatility of the returns. We observe that the volatility of the exchange rate returns closely follow each other, which is expected considering earlier observations of highly correlated returns. We do not observe the same for the salmon volatility, where the the spot series centers around a level that is approximately twice as high as for the forward, which is expected given the descriptive statistics in Table 4.1. Furthermore, we see that salmon returns in general are far more volatile than the exchange rate returns, confirming the need for efficient risk management practices in the salmon farming industry.

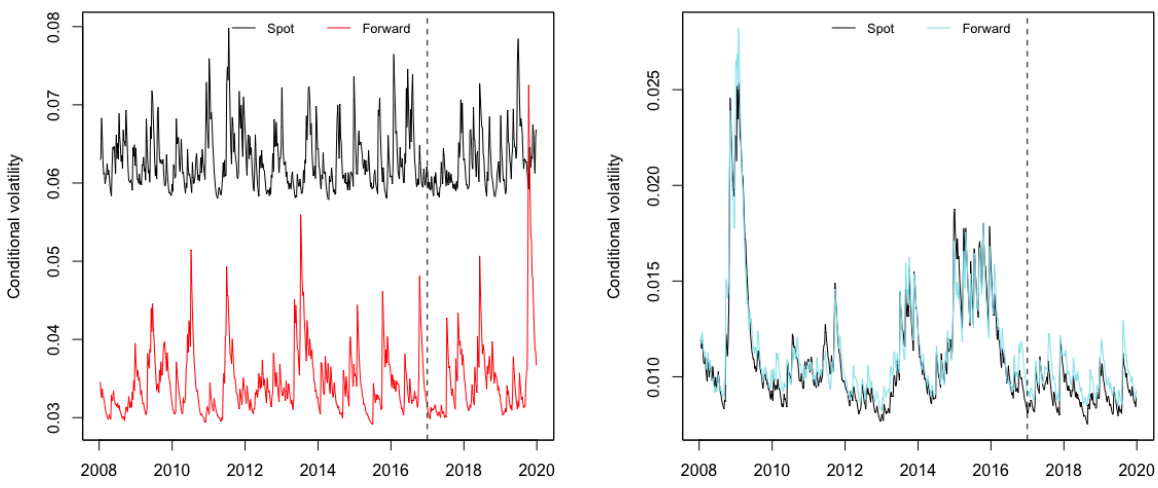

(a) Salmon **(b)** Currency

Figure 4.3: GARCH(1,1) - modeled conditional volatility of salmon and currency.
Note: Dashed line divides in- and out-of-sample data.

We present the estimated DCC(1,1) models in Table 4.5. This table shows several interesting results. First, we observe that the α - parameter is strongly insignificant in the salmon single model. We cannot reject the null that $\alpha = 0$, as the magnitude of the p-value indicates strong evidence for the null hypothesis. This result is reasonable when looking at the modeled correlation between the salmon spot and forward in Figure 4.4a. A small α indicates that short term shocks have little to no effect on the correlation. The size of α determines the effect of past innovations, and therefore the short-term persistence in the series. Intuitively, the salmon market is still immature and would most likely react slowly to changes or shocks. Consequently, we observe the correlation to move in a periodical manner in Figure 4.4a.

Table 4.5: Estimated DCC(1,1) models for single hedges and multi hedge.

	Single salmon		**Single currency**		**Multi hedge**	
	Est.	p-value	Est.	p-value	Est.	p-value
α	0.0041	0.8124	0.0940	0.0364	0.0431	0.0027
β	0.9451	0.0000	0.1295	0.6251	0.0000	1.0000

Notes: Model parameter estimates highlighted in grey are not statistically significant at the 5% level.

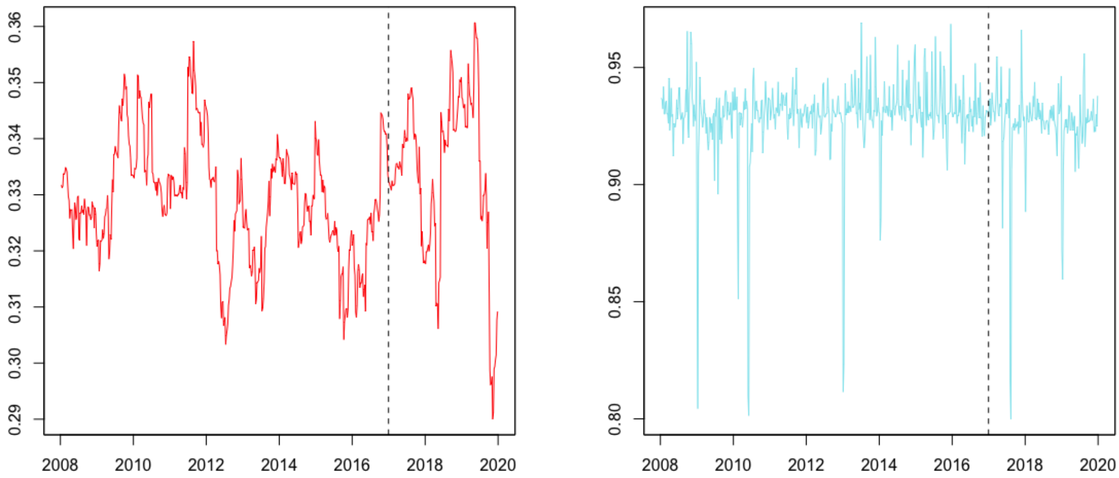

(a) Salmon

(b) Currency

Figure 4.4: DCC(1,1) - modeled conditional correlation for single hedges of salmon and currency. *Note:* Dashed line divides in- and out-of-sample data.

Second, we observe that the β - parameter is not statistically significant in the currency single model. In other words, we cannot reject the null that $\beta = 0$. The β indicates the effect of lagged conditional variance on the correlation, and is therefore a measure of the long-run persistence of clustering in the series. If $\alpha + \beta$ is high the correlation will decay slowly. As β is low, or very likely 0, correlation will decay quickly after shocks. This is confirmed by Figure 4.4b, which plots the modeled correlation between the exchange rate spot and forward. Considering that the FX market is mature, it is reasonable to believe that it reacts and adjusts quickly, meaning shocks and changes do not persist.

Lastly, we observe that the β - parameter equals 0 for the multi hedge. This indicates that the lagged conditional variances do not affect the conditional correlations between the salmon and exchange rate returns and, hence, long term effects have no impact.

Chapter 5

Results

In this chapter, we provide answers to three questions. First, do dynamic hedging strategies, modeled with the GARCH framework, outperform the naïve hedge? Second, does the multi-product hedge that takes dependencies between salmon and currency into account, outperform the single hedge? Third, do the threshold strategies, that aim to take advantage of volatility clustering, yield competitive risk management opportunities for salmon producers compared to more traditional strategies?

Our main findings can be summarized as follows. First, we find that the dynamic hedging strategies outperform the naïve hedge in terms of mean return and cost, but not in terms of variance reduction. Second, we find that the multi hedge performs similarly to the single hedge in-sample, but outperforms the single hedge out-of-sample, when prices are depreciating. Third, we find that the threshold strategies perform similarly to the traditional strategies in-sample, while substantially outperforming them out-of-sample.

The rest of this chapter is organized as follows. Section 5.1 presents the hedge ratios that different strategies yield. In Section 5.2 we look at the hedging performance of the strategies in terms of mean return and variance reduction for the four-week horizon. Then we conduct a sensitivity analysis of the hedging results for longer horizons in Section 5.3. The transaction costs for the different strategies are considered and analysed in Section 5.4, followed by a discussion of the robustness of the results in Section 5.5.

5.1 Hedge ratios

We first compare the dynamic hedge ratios for the single and multi hedge in Section 5.1.1. While the single hedge only considers the conditional volatility of the underlying product, the multi hedge exploits potential dependencies between the salmon and exchange rate returns. Then we assess the hedge ratios for the threshold strategies in Section 5.1.2.

5.1.1 DCC-GARCH

In order to see if multi-product hedging is effective we compare the single hedge and multi hedge optimal hedge ratios (OHR). In Figure 5.1 we see the dynamic OHRs for salmon and currency, with the left panel showing the OHR for the single hedge and the right panel for multi hedge. Comparing the OHR for salmon for the single hedge and the multi hedge, we see that they are similar. This is also reflected in the descriptive statistics in Table 5.1 which shows the mean, standard deviation (SD), minimum and maximum for the OHRs in the single and multi hedge. Here, we see that the mean OHRs for salmon in the single hedge (0.6025) and the multi hedge (0.5889) are close. The SD for multi hedge (0.0822) is a little higher than the SD for the single hedge (0.0695). The higher volatility for the multi hedge is also confirmed by the higher maximum (0.8863) and lower minimum (0.2244) OHRs. However, the small deviations between the OHRs for the single hedge and multi hedge indicate that the salmon price is close to optimally hedged in the single hedge, when currency is not taken into account. Hence, the added benefit in the multi hedge should come from the adjusted positions in the exchange rate derivatives. This is in line with the findings of Benninga and Eldor (1985), who find that the commodity hedge is independent of the FX market, while the OHR for the exchange rate depends on the commodity hedge.

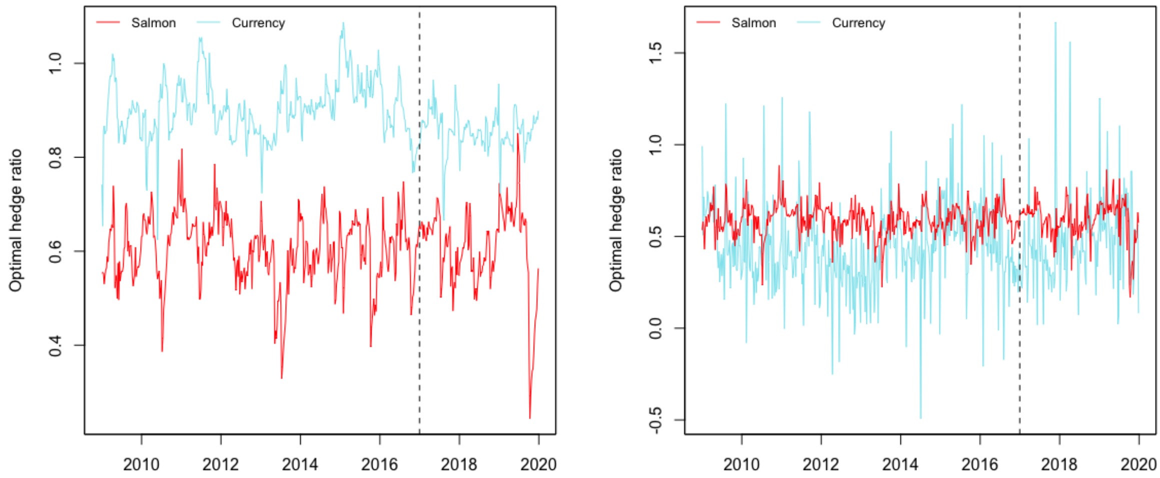

(a) Single hedge **(b)** Multi hedge

Figure 5.1: Hedge ratios for dynamic strategies using DCC-GARCH.
Note: Dashed line divides in- and out-of-sample data.

Looking at Figure 5.1 of the OHR for the exchange rates in the single hedge and multi hedge, we observe a substantial difference between the two. In Table 5.1, the single hedge has a mean OHR for the exchange rates (0.9034), that is substantially higher than for the multi hedge (0.4400). While the mean OHR is lower for the multi hedge, the SD is higher (0.2285) than the single hedge (0.0631). The higher volatility is also reflected in the higher maximum (1.2576) and lower minimum (-0.4921) for the multi hedge.

The changes in the OHRs for exchange rate in multi hedge compared to single hedge indicate that the multi hedge effectively captures dependencies between the currency and salmon. We also observe that the multi hedge uses the exchange rate futures more actively to offset unfavorable price movements in the salmon market. Intuitively, this can be explained by the efficiency in the FX markets (Rapp and Sharma, 1999) and the inefficiency in the salmon market (Chen and Scholtens, 2019). This is evident when looking at the correlation between the spot and forward prices in the respective markets. The exchange rates exhibit a correlation of 0.92, while for salmon this correlation equals 0.31. The higher correlation in the FX market explains the larger positions in currency derivatives compared to the

Table 5.1: In-sample hedge ratios for dynamic strategies.

	Single	Multi
Salmon		
Mean	0.6025	0.5889
SD	0.0695	0.0822
Min	0.3496	0.2244
Max	0.8218	0.8863
Currency		
Mean	0.9034	0.4400
SD	0.0631	0.2285
Max	0.6245	-0.4921
Min	1.0874	1.2576

positions in salmon derivatives, as hedging is more effective when the spot and forward are more correlated.

5.1.2 Threshold strategies

To investigate the performance of the threshold strategies, we describe the hedge ratios for the naïve threshold hedge and the dynamic threshold hedge. Figure 5.2 illustrates the obtained hedge ratios. The threshold strategies exploit volatility clustering by hedging when recent returns reach the threshold, and otherwise holding an unhedged position.

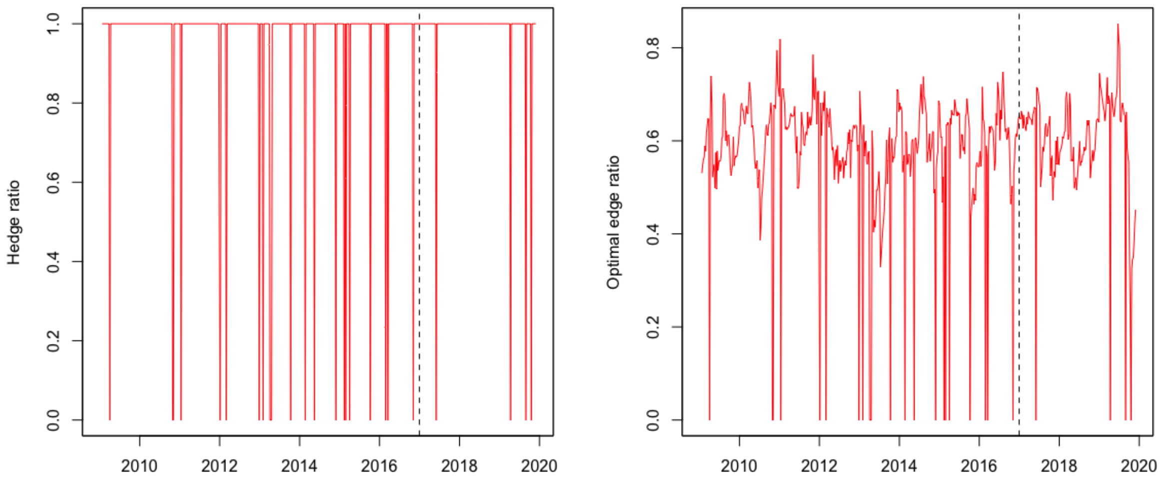

(a) Naïve threshold hedge

(b) Dynamic threshold hedge

Figure 5.2: Hedge ratios for salmon for the threshold strategies.
Note: Dashed line divides in- and out-of-sample data.

The hedge ratio for the naïve threshold hedge can be seen in Figure 5.2a and in Table 5.2. Here, we see that the strategy results in the second highest mean hedge ratio (0.9471) of all strategies, only beaten by the naïve hedge. For the dynamic threshold hedge, we expect a lower hedge ratio compared to the single hedge, since a hedge is only initiated when the threshold is reached. Figure 5.2b and Table 5.2 confirm this, as the strategy has a mean hedge ratio (0.5706) which is slightly lower than the single hedge (0.6025). This indicates that the threshold is often reached due to frequent volatility clustering, which results in holding the single hedge position most of the time. The large changes between a hedged and unhedged position results in large deviations in hedge ratio. The dynamic threshold hedge SD (0.1512) is therefore substantially higher than for the single hedge (0.0695) and multi hedge (0.0822).

Table 5.2: In-sample hedge ratios for threshold strategies.

	Naïve	Dynamic
Salmon		
Mean	0.9471	0.5706
SD	0.2241	0.1512
Min	0.0000	0.0000
Max	1.0000	0.8218

The descriptive statistics for the hedge ratios in the out-of-sample period can be seen in Table B.1 and Table B.2. The slightly higher mean hedge ratio for the out-of-sample period for the naïve threshold hedge compared to in-sample period indicates more volatility clustering out-of-sample.

These results indicate that the threshold strategies hold similar positions to the naïve hedge and single hedge most of time. In Section 5.2 we will investigate if threshold strategies adds value in terms of return and variance reduction.

5.2 Hedging results

In this section, we further explore the different strategies by comparing how they perform in terms of portfolio return and variance reduction. The portfolio return is based on Equation 3.1, using the data and hedge ratios presented earlier. The variance reduction is evaluated using the hedge effectiveness (HE) and expected shortfall measures described in Section 3.4.

The summary of the hedging results for a four-week horizon is presented in Table 5.3. As described in Chapter 4, we divide our data into two sub-samples, in- and out-of-sample. The in-sample mean return and HE explains how well the strategies adapt to the training set and how well they perform on known data, while the out-of-sample results are used to evaluate the forecasting performance. The out-of-sample results could therefore be of more interest to a salmon producer as the strategies are tested on unknown data. The return paths for the six strategies are presented in Figure B.1. First, we look at the in-sample results, then the out-of-sample results, before comparing the two.

Table 5.3: Hedging results for the four-week horizon.

	Unhedged	Naïve hedge	Single hedge	Multi hedge	Naïve threshold	Dyn. threshold
In-sample						
Return outcome						
Mean return	796	24	318	325	87	358
Min return	-34890	-30392	-32235	-32003	-30248	-32099
Max return	24905	27168	26384	26034	26611	25936
Mean transaction cost	-	263	158	155	249	150
Variance outcome						
Standard deviation	8360	6963	7061	7130	7084	7236
Hedge effectiveness	-	30.63%	28.65%	27.26%	28.20%	25.09%
ES 5% | Reduction	17943	14.13%	12.54%	11.73%	13.28%	10.94%
ES 10% | Reduction	14296	12.86%	13.16%	11.92%	11.77%	10.84%
Out-of-sample						
Return outcome						
Mean return	-181	-211	-108	20	152	187
Min return	-33681	-25068	-22077	-23837	-24649	-22630
Max return	33680	26209	28026	28493	27485	29083
Mean transaction cost	-	263	160	156	256	156
Variance outcome						
Standard Deviation	13241	9050	10366	10440	9243	10531
Hedge effectiveness	-	53.28%	38.70%	37.83%	51.27%	36.73%
ES 5% | Reduction	24184	30.80%	20.75%	20.06%	30.00%	21.05%
ES 10% | Reduction	16757	13.76%	1.68%	1.87%	11.89%	0.42%

Note: Returns denoted in 1000s NOK.

In-sample

Comparing the strategies, we see that the unhedged portfolio is the most volatile. This is not surprising as substantial portfolio volatility is the primary motivation for why producers want to engage in hedging. We also observe that the unhedged portfolio delivers the highest mean return (796) over the in-sample period, which is expected as the salmon prices on average appreciate over this period. Moreover, we observe that the naïve hedge yields the highest HE (30.63%) and lowest mean return (24), of all strategies. This is in line with Asche et al. (2016), who find that the naïve hedge outperforms other hedging strategies in terms of variance reduction. In addition, we note that this strategy also has the highest mean transaction cost (263), which is expected considering a full futures position is held at all times. Both the single hedge (318) and the multi hedge (325) yield higher mean returns than the naïve hedge, while still achieving fairly high HE of 28.75% for the former and 27.26% for the latter. The transaction costs for the two strategies are also substantially reduced (158 and 155) compared to the naïve hedge. Looking at the threshold strategies, we see that the naïve threshold hedge is close to the naïve hedge, and the dynamic threshold hedge is close to the single and multi hedge. The naïve threshold hedge yields a lower mean return (87) compared to the dynamic threshold hedge (358), but achieves a higher HE (28.20% and 25.16%). The naïve threshold hedge costs less (249) than the naïve hedge, while the dynamic threshold hedge (150) has the lowest transaction costs of all strategies.

In addition, we use expected shortfall (ES) to evaluate tail risk of the proposed strategies. We expect the unhedged portfolio to have the highest ES as hedging would prevent several of the worst case scenarios. As seen from Table 5.3 all hedging strategies reduce the ES compared to the unhedged portfolio, confirming our expectation. At the 5% level, the naïve hedge (14.13%) yields the best reduction in ES, while the single hedge (13.16%) yields the best results at the 10% level. The 10% results indicate that the single hedge performs better than the other strategies at reducing the average loss for a bigger fraction of the adverse scenarios. However, the 5% results indicate that the naïve hedge and naïve threshold hedge are better at hedging the few very extreme scenarios.

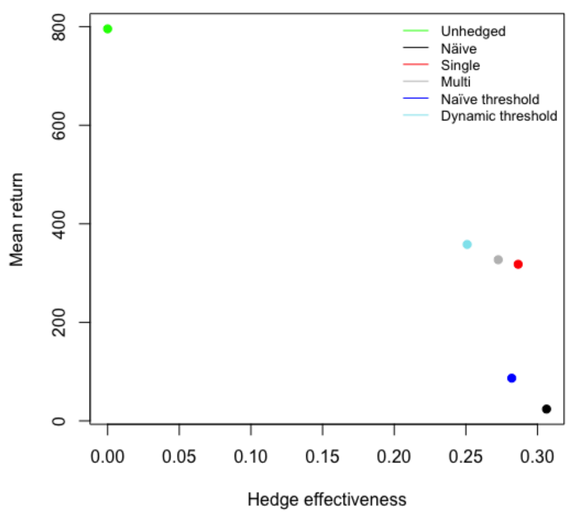

Figure 5.3: In-sample mean return and hedge effectiveness for all strategies.

Figure 5.3 plots the mean return against the hedge effectiveness. We use this to summarize the in-sample findings. The best performing strategy would be in the upper right quadrant, as it would achieve the best combination of HE and mean return. First, we observe that the single and multi hedge yield considerably higher mean return and lower transaction costs, while performing similarly in terms of HE. Second, there are no substantial differences between the single and multi hedge. Third, we observe that the dynamic threshold hedge yields a slightly higher mean return at the cost of HE compared to the single and multi hedge. The same applies to the naïve threshold hedge, which also yields a higher mean return and lower HE compared to the naïve hedge.

Out-of-sample

In this section, we consider the out-of-sample hedging results, which occurs during a period of depreciating salmon prices. An overview of the mean return plotted against the HE for the sample is presented in Figure 5.4.

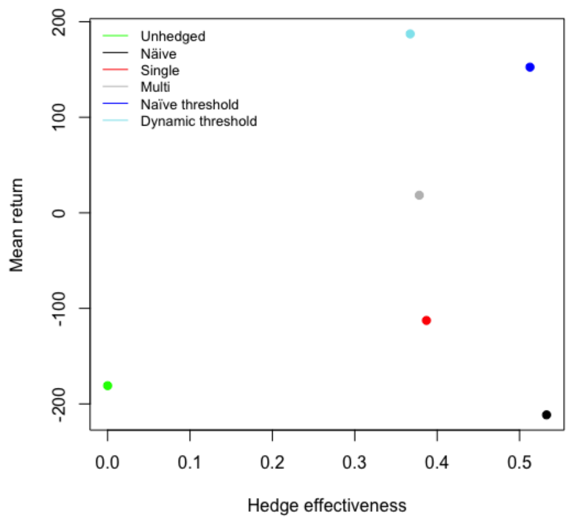

Figure 5.4: Out-of-sample mean return and hedge effectiveness for all strategies.

First, we consider the HE of our strategies. Several of the results coincide with the in-sample observations. The naïve hedge still delivers the highest HE (53.28%), while the single hedge (38.70%) and multi hedge (37.83%) deliver quite similar results to each other. From Table 5.3 we see that the naïve threshold hedge (51.27%) yields a higher HE compared to the findings of Asche et al. (2016), which found that salmon derivatives reduce price risk with 30-40%. A high HE is also seen for the dynamic threshold hedge (36.73%), which achieves approximately similar HE to the single and multi hedge.

We observe that the strategies are not able to outperform the naïve hedge in terms of HE. However, regarding mean return and transaction costs, the results differ substantially from the in-sample results. The mean return for the unhedged portfolio (-181) is negative. Moreover, both dynamic strategies yield higher mean return than the naïve hedge (-211), which has the highest transaction costs (263). The difference between the mean return of single (-108) and multi hedge (20) is also noticeable, while the transaction costs are fairly similar. Furthermore, we observe that the threshold strategies yield superior results out-of-sample, with a mean return of 152 for the naïve threshold hedge and 187 for the dynamic threshold hedge. The transaction costs are somewhat lower (263 and 156) compared to the respective naïve and single hedge.

Lastly, we examine the expected shortfall (ES). It can be seen that the ES for the dynamic threshold hedge (21.05%) is fairly close to the single hedge (20.75%) and multi hedge (20.06%), while the naïve threshold hedge (30.00%) is close to the naïve hedge (30.80%). This means that the threshold strategies do not remove down-side tail risk as effectively as the traditional strategies. However, since the mean return is substantially higher, the strategies seems to capture the positive spikes, which are necessary if the producers are to deliver positive returns when prices depreciate.

To summarize, the dynamic strategies perform substantially better in terms of mean return and transaction costs than the naïve hedge, while they maintain a relatively high HE. Moreover, we observe that by utilizing a multi hedge we gain a considerably higher mean return and slightly lower transaction costs compared to the single hedge, while performing similarly in terms of HE. Last, the threshold strategies yield superior results with regards to both mean return and HE.

In-sample versus out-of-sample

In this section, we compare the in- and out-of sample results. First, we look at the distributions of portfolio returns for the different strategies. Then we look at the differences of the in- and out-of-sample results.

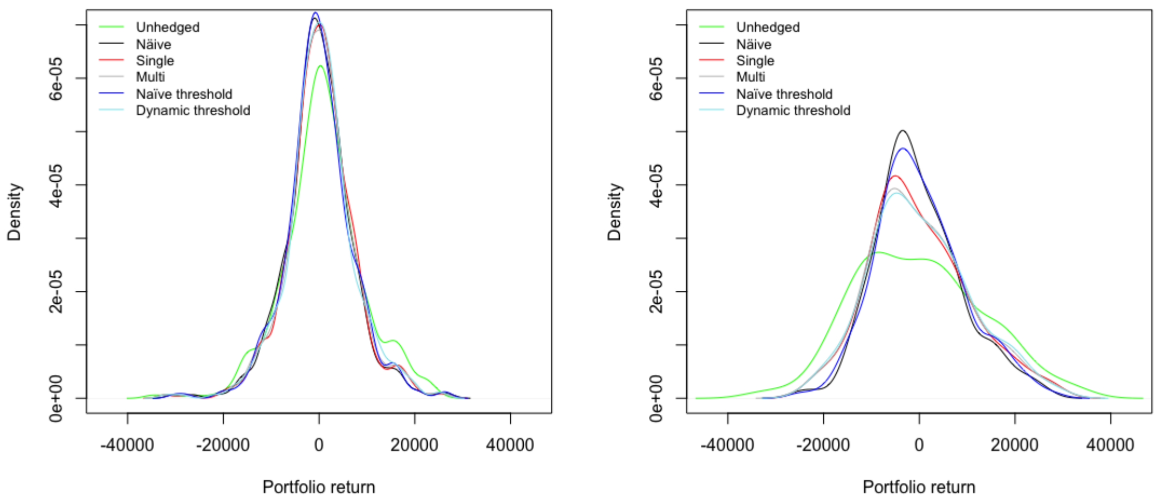

(a) In-sample **(b)** Out-of-sample

Figure 5.5: Density of portfolio returns for the in- and out-of-sample periods.

Figure 5.5 shows the distribution of the portfolio returns for the strategies in- and out-of-sample. For the in-sample period, we see that the unhedged portfolio has a fatter right tail and fewer observations around 0, than the other strategies. Furthermore, the distributions of all hedging strategies are fairly close to each other. For the out-of-sample period, we observe a more substantial difference between the strategies. We note that the unhedged strategy again has fatter tails compared to the other strategies.

This indicates that positive returns are offset by negative returns. The strategies with the highest HE have shorter tails and have a higher concentration of observations around 0. We see that the naïve hedge yields the lowest variance in portfolio return but as a result it also has fewer positive observations than other strategies. The naïve threshold hedge has a short tail in the left part of Figure 5.5b while having a longer tail on the right side, resulting in a high HE and mean return. This confirms our findings in Table 5.3.

When comparing the in- and out-of-sample hedging results from previous sections, we notice several interesting observations. First, it is clear that the naïve hedge performs best in terms of HE. However, with high HE, the naïve hedge removes most of the upside in both appreciating and depreciating times, leading to its low mean return as evidenced by both sample periods. Second, we observe that dynamic strategies increase mean return and reduce the transaction costs without sacrificing too much HE, both in- and out-of-sample compared to the naïve hedge. In addition, expanding the single hedge to a multi hedge may substantially increase portfolio returns as proven by our out-of-sample results. Third, we see that the dynamic threshold hedge yields the highest mean return in both samples. By utilizing the predictive powers provided by the volatility clustering, it is able to capture more of the positive returns compared to the other strategies. A higher mean return is also seen when we compare the naïve threshold hedge to the naïve hedge.

The differences between the in- and out-of-sample results can be attributed to several factors, out of which, the price dynamics of the underlying periods seems to be the most important. When comparing to the unhedged results, hedging is more beneficial for the depreciating out-of-sample period than to the appreciating in-sample-period.

Our findings indicate that the preferred hedging strategy would depend on the risk-preference of the producer. Current industry hedging practices are simple as they are usually static.[1] By employing dynamic strategies, a producer could obtain higher mean returns, while still reducing variance and transaction costs effectively. Even higher mean returns can be obtained through the dynamic threshold hedge as seen in both samples. However, the complexity of modeling the conditional volatility is a barrier. The naïve threshold hedge is easier to implement, and thus extend the producers repertoire of available strategies. As seen, this strategy yields a higher mean return than the naïve hedge, while preserving a similar HE.

5.3 Sensitivity to hedging horizon

In this section, we analyze our hedging results with respect to different horizons. We extend the four-week horizon by rolling the front month futures contract.[2] The in- and out-of-sample periods are combined to include effects of both appreciating and depreciating prices. Figure 5.6 depicts the HE and mean return for the one- to ten-week horizon for all strategies. First, we look at the HE, then the mean return.

Looking at Figure 5.6a, we observe that HE increases with longer horizons for all strategies. This

[1] This insight was revealed in a phone interview with an industry representative.

[2] Futures can be rolled from the front-month contract close to expiration to another contract in a further-out month, this way the contract does not need to be settled and the trader avoids the costs and obligations associated with the mentioned settlement.

finding is in line with Haarstad et al. (2021) and Bloznelis (2018). For all horizons, we see that the naïve hedge yields the highest HE. When comparing the single and multi hedge, we observe that they closely follow each other until the four-week horizon, whereas afterwards the gap between them slowly increases. However, we note that the divergence between them is within a few percentage points. When considering the threshold strategies, we observe that the naïve threshold hedge is right below the naïve hedge when it comes to HE. The dynamic threshold hedge yields the lowest HE, just below the single hedge.

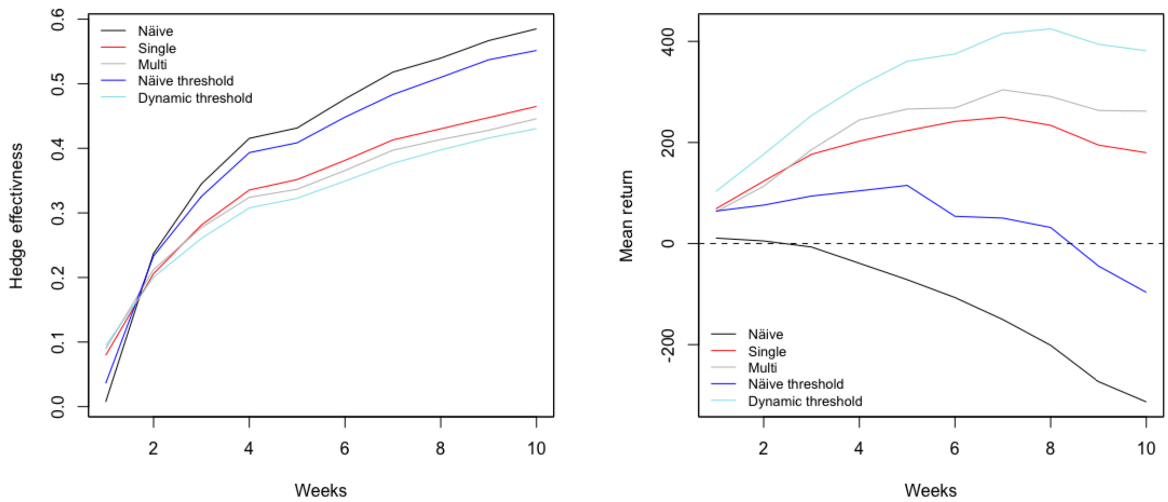

(a) Hedge effectiveness **(b)** Mean return

Figure 5.6: Mean return and hedge effectiveness for full sample period.

From Figure 5.6b, we see substantial differences between the strategies. First, we observe that the naïve hedge has a positive mean return from one to three weeks, but negative for longer horizons. It steadily decreases for longer horizons, in line with the increasing HE. This is a result of over-hedging. Second, the multi hedge yields a higher mean return than the single hedge, especially for longer horizons. Both strategies increase until week seven, then they start to decline. Last, the two threshold strategies yield very different results compared to each other. The dynamic threshold hedge performs best of all in terms of mean return for all horizons, while the naïve threshold hedge has a mean return in between the single and naïve hedge. Moreover, the mean return appreciates until week five for the naïve threshold hedge, whereas the dynamic threshold hedge appreciates until week eight.

Several of the findings from the sensitivity analysis coincide with earlier findings for the four-week horizon. First, the naïve hedge outperforms the dynamic strategies in terms of HE, while the dynamic strategies yield substantially better mean return for all horizons. Second, the single hedge and multi hedge perform similarly with respect to HE, while the multi hedge achieves higher mean return for ever longer horizons. Third, we again observe the potential of the threshold strategies. The naïve threshold hedge performs similar to the naïve hedge with regards to HE, while it yields noticeably higher mean return for all horizons. For the dynamic threshold hedge we observe that it performs similar to the single and multi hedge in terms of HE, but yields superior mean return to all strategies.

This sensitivity analysis also provides other useful insights. Most notably, the hedge effectiveness seems to increase for longer horizons for all strategies. Therefore it would be beneficial for risk averse salmon producers to roll the futures contracts over longer horizons. For most strategies mean return increases for shorter horizons and decreases for longer horizons. Thus, a longer hedging horizon is not necessarily preferred when we take into account return considerations. The single, multi and dynamic threshold hedge reach their maximums around week 7-8, which in practice means rolling the front month contract once. To summarize, there is potential value for salmon producers, both in terms of return and risk reduction, by considering longer horizons.

5.4 Cost of hedging

As we observe that superior variance reduction comes at the cost of mean return, it is important to compare the strategies in terms of costs. This is useful as the size of hedging positions, and therefore the costs, differ from strategy to strategy. Moreover, the costs could indicate the viability of the hedging strategies for salmon producers, as high costs are likely to deter producers from hedging. First, we look at the transaction costs by adjusting the mean returns of the different strategies. Second, we examine the cost of variance reduction.

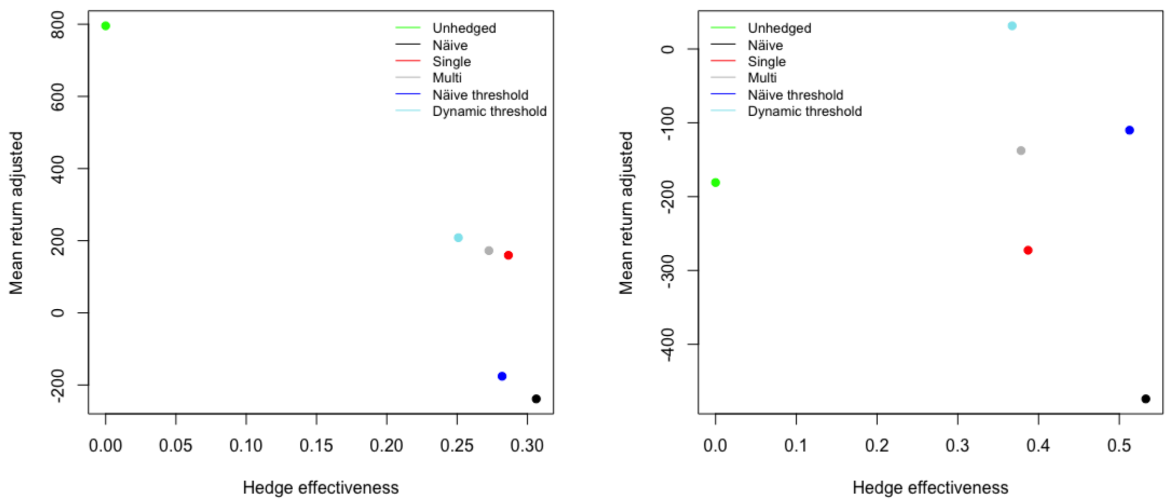

(a) In-sample **(b)** Out-of-sample

Figure 5.7: In- and out-of-sample mean return adjusted for transaction costs.

Figure 5.7 depicts the mean return adjusted for the transaction costs of the different strategies both in- and out-of-sample. For the in-sample period, we observe a similar picture as we did without adjusting for transaction costs, apart from the fact that the naïve and the naïve threshold hedge take on negative mean returns due to high transaction costs. For the out-of-sample period, we observe that the single hedge and the naïve hedge perform worse than the unhedged, meaning that the transaction costs offset the benefits in terms of return for these strategies. Furthermore, we see that the gap between the dynamic threshold hedge and the others has widened. Especially noticeable is the return difference between the dynamic and naïve threshold hedge. The extra reduction in variance gained from the naïve threshold hedge comes with a high cost, resulting in a negative adjusted mean return. The benefit of the dynamic threshold hedge is also evident in Figure 5.8, which plots the out-of-sample adjusted mean return against hedging horizon. The dynamic threshold hedge is the only strategy with positive mean return during the period. Therefore, when solely considering return, the dynamic threshold hedge is the obvious choice for a salmon producer. Additionally, we note that when the horizon exceeds four weeks, the naïve threshold and the multi hedge also perform better than the unhedged alternative in terms of mean return.

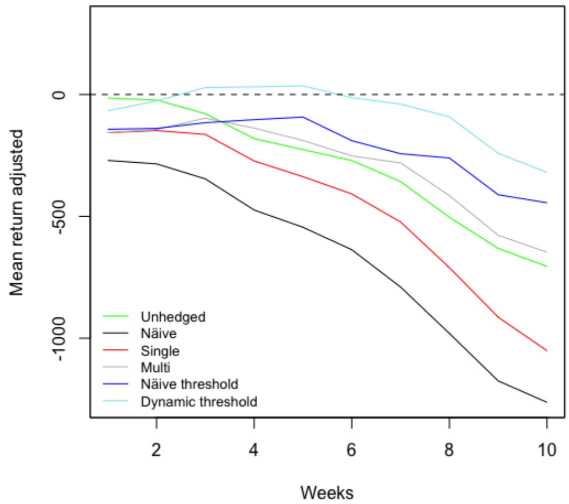

Figure 5.8: Mean return adjusted for transaction costs for the one- to ten-week horizon.

In order to evaluate the costs of variance reduction, we look at the trade-off between risk and return. This is done by analyzing the cost of variance reduction (CVR) measure proposed by Haarstad et al. (2021), which is equal to the cost of hedging, divided by the hedge effectiveness. The cost of hedging is calculated by subtracting the adjusted mean return of the underlying hedging strategy from the unhedged return, which we use as a reference. As this measure essentially is the cost of hedging per percentage point of variance reduction, smaller values are favorable.

Table 5.4 presents the cost of hedging and CVR for the hedging strategies both in- and out-of-sample with the four-week horizon. In general, we observe for the in-sample period that the simplest strategies, the naïve and the naïve threshold hedge, perform worse than the more advanced strategies. This is mainly due to the large cost of hedging, which results from high transaction costs and low mean return. This implies that the variance reduction gained from the simpler strategies are more costly than for the more advanced hedging strategies. Therefore, when employing the CVR measure, more advanced strategies perform better. This is also confirmed by the out-of-sample results, where the simplest strategy, the naïve hedge, performs the absolute worst.

Table 5.4: Summary of mean return adjusted, cost of hedging, hedge effectiveness and cost of variance reduction for the four-week horizon.

Hedge effectiveness	Unhedged	Naïve hedge	Single hedge	Multi hedge	Naïve threshold	Dyn. threshold
In-sample						
Mean return adjusted	796	-239	160	172	-162	208
Cost of hedging	0	1034	636	623	958	587
Hedge eff. (%)	-	30.63%	28.65%	27.26%	28.20%	25.09%
CVR	-	33.77	22.20	22.86	33.96	23.41
Out-of-sample						
Mean return adjusted	-181	-474	-273	-138	-103	31
Cost of hedging	0	293	92	-43	-78	-212
Hedge eff. (%)	-	53.28%	38.70%	37.83%	51.27%	36.73%
CVR	-	5.50	2.37	-1.14	-1.52	-5.78

Note: Returns denoted in 1000s NOK.

Looking more closely at the in-sample results in Table 5.4, we observe that the strategy with the lowest CVR is the single hedge (22.20), slightly ahead of the multi hedge (22.86). This is mainly due to the single hedge having a slightly higher HE and lower mean return, as discussed earlier. In terms of CVR for the in-sample period, there is no added benefit of taking the dependencies between salmon prices and exchange rates into account.

The CVR of the dynamic threshold hedge (23.41) is close to the single and multi hedge. The adjusted mean return is substantially higher than for the other hedging strategies, but the relatively low HE (25.09%) decrease the CVR. However, this indicates that the dynamic threshold hedge is worth considering for a producer that wants to reduce the variance while maintaining a higher mean return. We also observe that the naïve threshold hedge has the highest CVR (33.96), close to the naïve hedge (33.77). With ever longer horizons, we expect the naïve threshold hedge to yield better CVR results than the naïve hedge due to higher mean return, as seen in Figure 5.8.

For the out-of-sample period, the strategies achieving the lowest CVR are different. This is mainly due to the increased benefit of hedging when prices depreciate. The multi hedge (-1.14) achieves a negative CVR value, while the single hedge now has the second highest CVR (2.37). The negative CVR indicates that the hedged position yields a higher mean return adjusted for transaction costs compared to the unhedged position. This illustrates the value of using a multi-product hedge while prices are depreciating. In stark contrast to the in-sample period, the out-of-sample results indicate that taking dependencies between salmon prices and exchange rates into account is beneficial.

The dynamic threshold hedge has by far the lowest CVR (-5.78), as it manages to maintain a relatively high HE (36.73%) while also having the lowest cost of hedging. The naïve threshold hedge has a negative CVR (-1.52), which is considerably lower than the naïve hedge (5.50). Again, it can be seen that the threshold strategies achieve higher mean returns with lower transaction costs while having relatively high HE.

There are several key takeaways from this section. First, we observe that the dynamic strategies perform substantially better than the naïve hedge using the CVR measure. Second, the CVR results indicate that the multi hedge performs almost as well as the single hedge during the in-sample period while achieving a substantially better result out-of-sample. Third, the dynamic threshold hedge yields superior results during periods of depreciating prices, while in periods of appreciating prices, it performs similarly to the dynamic strategies. Moreover, the results indicate that the naïve threshold hedge could be considered as a direct substitute to the naïve hedge. In addition to being easy to implement in practice, they both perform similarly in-sample, while the naïve threshold hedge performs substantially better out-of-sample.

5.5 Robustness tests

To evaluate the robustness of our findings, we investigate the sensitivity of our results with respect to the choice of the in- and out-of-sample periods. We do this by dividing the full sample into smaller five-year sub-samples. The sub-samples are again divided into in-sample periods of four years and out-of-sample periods of one year.[3] The two measurements we evaluate in this section are the hedge effectiveness and mean return given a four-week horizon. The robustness results are presented in Table B.3 and Table B.4. The average of the results are summarized in Figure 5.9 where the HE is plotted against the mean return

[3] The sub-sample period is from 2009-2013 until 2013-2017.

for the in- and out-of-sample periods. First, we look at the in-sample robustness, then the out-of-sample robustness.

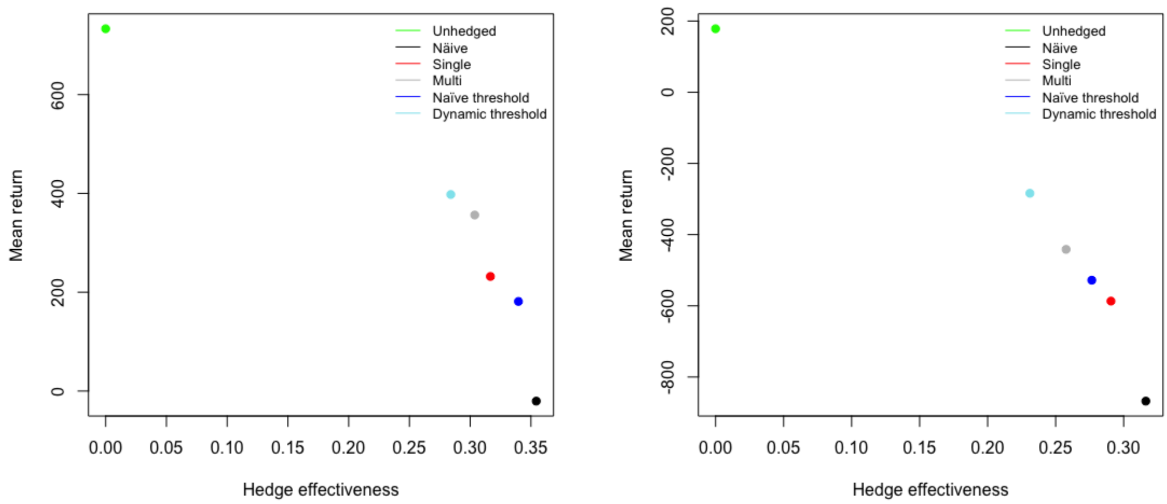

(a) In-sample **(b)** Out-of-sample

Figure 5.9: Average of mean return and hedge effectiveness for all in- and out-of-sample periods for the four-week horizon.

Looking at Figure 5.9a of the in-sample results, we observe a fairly linear relationship where the strategies with the highest HE have the lowest mean return and opposite. Thus, we observe a more prominent risk-return trade-off than what earlier results indicate. However, in general we observe that the original in-sample findings are confirmed. First, we see that the single and multi hedge yield substantially higher mean return than the naïve hedge, with somewhat lower HE. This confirms that using dynamic strategies is valuable for producers that want risk reduction without sacrificing too much return. Second, Figure 5.9a shows that the multi hedge achieves a higher mean return than the single hedge, while it achieves comparable HE. This confirms that it is beneficial to utilize a joint salmon price and currency hedging framework. Third, we observe that the threshold strategies perform in line with earlier findings. The naïve threshold hedge performs similarly in terms of HE to the naïve hedge while having a substantially higher mean return. The dynamic threshold hedge yields the highest mean return of all strategies while having an HE close to the multi hedge, as observed earlier. Overall, our original in-sample findings seems to be robust. This could be because the prices in the robustness test on average appreciate, similarly to the original in-sample period.

The results for out-of-sample data presented in Figure 5.9b indicate a close to linear risk-return relationship. This was not observed in the original out-of-sample data due to different price dynamics in the underlying periods. Earlier we noted the depreciating characteristics of the salmon prices for the original out-of-sample data, while for the periods used in the robustness test we observe appreciating prices. Therefore, we cannot directly compare these results. However, the out-of-sample robustness test results can be compared to the original in-sample results as the data share the same appreciating characteristics. Comparing Figure 5.9a and Figure 5.9b we observe similar features, and thus, the robustness of the original in-sample findings is further strengthened.

Chapter 6

Conclusion

In this thesis, we address two of the most important risks Norwegian salmon producers face, salmon price and currency fluctuations. As significant amounts of salmon are exported internationally, both these risks substantially impact producers' competitiveness and profitability. Therefore, efficient risk management practices are of great importance. The objective of our thesis was to investigate if salmon producers can improve current risk management practices by utilizing more advanced strategies for hedging both price and currency risk.

In what follows we summarize our main findings. First, we find that dynamic hedging strategies increase returns while performing similarly in terms of risk reduction, compared to the naïve hedge. The performance of the dynamic strategies is further substantiated when transaction costs are accounted for, as these are lower than for the naïve hedge. Thus, we conclude there is value in using dynamic hedging strategies for salmon producers. Second, we find that utilizing dependencies between the salmon and currency markets through a state-of-the-art multi-product hedging framework is beneficial. Compared to hedging the products independently, the multi-product hedge performs similarly in terms of risk reduction, but yields higher returns and lower transaction costs. Third, we find that novel threshold strategies yield higher returns than more traditional strategies, and performs similarly in terms of risk reduction. Thus, we find value in considering volatility clustering effects for hedging purposes. Fourth, we consider the costs of hedging compared to variance reduction, as it is an indication of the viability of different hedging strategies for salmon producers. In this regard, we find that more advanced hedging strategies perform better than the naïve hedge. Fifth, we find that the hedging horizon greatly affects the hedging results. Longer horizons yield increased risk reduction, but comes with a risk-return trade-off as returns eventually peak and start to decrease. Sixth, we find that the development of the salmon prices heavily influence the hedging performance. All considered strategies, except for the naïve hedge, improve both risk reduction and return when prices depreciate. The same is not observed when prices appreciate, indicating that hedging becomes more important when prices depreciate. Lastly, through these findings, we believe that salmon producers can improve current risk management practices by utilizing more advanced hedging strategies. By providing insights and practical steps on the implementation of dynamic and threshold strategies, that outperform the naïve hedge, this study is a valuable contribution to the industry. Moreover, it is also useful for other industries as the hedging strategies are applicable across commodities and currencies.

In what follows we propose interesting topics for further research. As we observe that more advanced hedging strategies often yield higher returns compared the traditional naïve hedge, a promising direction could be to investigate other hedging frameworks than the minimum-variance. Using a CVaR-framework that focuses on eliminating down-side risk and maximizing returns could yield interesting results. Moreover, further research could incorporate asymmetric properties in the modelling of the volatility in order to investigate if more of the positive shocks could be captured compared to the negative. Such asymmetry could be incorporated in the GARCH framework. Another direction that could be explored considers harvest rate. As described, our study assumes continuous harvest, which is consistent with current literature practice. In reality, the salmon producer has more flexibility when deciding on how much and when to harvest. Thus, our model could be further developed by incorporating varying harvest volumes. Lastly, a natural extension of our model is to include other currencies in order to investigate if further dependencies between markets could be utilized. For instance, the US market is important for Norwegian producers, and thus it would be interesting to include the US Dollar.

Bibliography

Acerbi, C., Nordio, C. & Sirtori, C. (2001). Expected shortfall as a tool for financial risk management. *Abaxbank working paper.*

Acerbi, C. & Tasche, D. (2002). On the coherence of expected shortfall. *Journal of Banking & Finance*, *26*(7), 1487–1503.

Allayannis, G. & Ofek, E. (2001). Exchange rate exposure, hedging and the use of foreign currency derivatives. *Journal of International Money and Finance*, *20*(2), 273–296.

Asche, F., Hansen, H., Tveteras, R. & Tveterås, S. (2009). The salmon disease crisis in chile. *Marine Resource Economics*, *24*(4), 405–411.

Asche, F., H.Roll, K., Sandevold, H. N., Sørvig, A. & Zhang, D. (2013). Salmon aquaculture: Larger companies and increased production. *Aquaculture Economics & Management*, *17*(3), 322–339.

Asche, F., Misund, B. & Oglend, A. (2016). Determinants of the atlantic salmon futures risk premium. *Journal of Commodity Markets*, *2*(1), 6–17.

Asche, F., Misund, B. & Oglend, A. (2018). The case and cause of salmon price volatility. *Marine Resource Economics*, *34*(1).

Asche, F., Misund, B. & Oglend, A. (2019). The case and cuase of salmon price volatility. *Marine Resource Economics*, *34*(1).

Bauwens, L., Laurent, S. & Rombouts, J. V. (2006). Multivariate garch models: A survey. *Journal of Applied Econometrics*, *21*(1).

Benninga, S. & Eldor, R. (1985). Optimal international heging in commodity and currency forward markets. *Journal of International Money and Finance*, *4*(4), 537–552.

Bergfjord, O. J. (2007). Is there a future for salmon futures? an analysis of the prospects of a potential futures market for salmon. *Aquaculture Economics Management*, *11*(2), 113–132.

Berk, J. & DeMarzo, P. (2014). *Corporate finance.* Pearson.

Bjørndal, T. & Tusvik, A. (2019). Economic analysis of land based farming of salmon. *Aquaculture Economics Management*, *23*(1), 449–475.

Bloznelis, D. (2016). Salmon price volatility: A weight-class-specific multivariate approach. *Aquaculture Economics Management*, *20*(1), 24–53.

Bloznelis, D. (2018). Hedging salmon price risk. *Aquaculture Economics Management*, *22*(2), 168–191.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, *31*(3), 307–327.

Bollerslev, T. (1990). Modelling the coherence in short-run nomial exchange rates: A multivariate generalized arch model. *The Review of Economics and Statistics*, *72*(3), 498–505.

Bollerslev, T., Engle, R. & Wooldridge, J. M. (1988). A capital asset pricing model with time-varying covariances. *Journal of Political Economy*, *96*(1), 116–131.

Brooks, C. (2014). *Introductory econometrics for finance.* Cambridge University Press.

Brooks, C., Henry, O. T. & Persand, G. (2002). The effect of asymmetires on optimal hedge ratios. *The Journal of Business*, *75*(2), 333–352.

Butterworth, D. & Holmes, P. (2001). The hedging effectiveness of stock index futures: Evidence for the ftse-100 and ftse-mid250 indexes traded in the uk. *Applied Financial Economics*, *11*(1), 57–68.

Caporin, M. & McAleer, M. (2009). *Do we really need both bekk and dcc? a tale of two covariance models.*

Chakraborty, A. & Barkoulas, J. T. (1999). Dynamic futures hedging in currency markets. *The European Journal of Finance*, *5*(4), 299–314.

Chang, C.-L., Gonzalez-Serrano, L. & Jimenzes-Martin, J.-A. (2013). Currency hedging strategies using dynamic miltivariate garch. *Mathematics and Computers in Simulation*, *94*, 164–182.

Chen, X. & Scholtens, B. (2019). The spot-forward relationship in the atlantic salmon market. *Fisheries Science and Aquaculture*, *27*(2), 142–151.

Creti, A., Joëts, M. & Mignon, V. (2013). On the links between stock and commodity markets volatility. *Energy Economics*, *37*(1), 16–28.

Ederington, L. H. (1979). The hedging performance of the new futures markets. *The Journal of Finance*, *34*(1), 157–170.

Engle, R. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, *20*(3), 339–350.

Engle, R. & Kroner, K. F. (1995). Multivariate simultaneous generalized arch. *Econometric Theory*, *11*(1), 122–150.

Fackler, P. L. & McNew, K. P. (1993). Multiproduct hedging: Theory, estimation, and an application. *Review of Agricultural Economics*, *15*(3), 521–535.

Fish Pool. (2020). *Fish pool abc.* Retrieved April 20, 2021, from https://fishpool.eu/wp-content/uploads/2014/08/Intro-EN-2020.pdf

Fisher, T. J. & Gallagher, C. M. (2012). New weighted portmanteau statistics for time series goodness of fit testing. *Journal of the American Statistical Association*, *107*(498), 777–787.

Grieg Seafood ASA. (2021). *Annual report 2020.* Retrieved May 1, 2021, from https://cdn.sanity.io/files/1gakia31/production/d2f76e0498525c2b0d07831dec33f4d2ed5547ed.pdf

Haarstad, A. H., Lavrutich, M., Strypet, K. & Strøm, E. (2021). Multi-commodity price risk hedging in the atlantic salmon farming industry. *Journal of Commodity Markets.*

Haigh, M. S. & Holt, M. T. (2002). Hedging foreign currency, freight, and commodity futures portfolios — a note. *The Journal of Futures Markets*, *22*(12), 1205–1221.

Hansen, P. R. & Lunde, A. (2005). A forecast comparison of volatility models: Does anything beat a garch(1,1)? *Journal of Applied Econometrics*, *20*(7).

Hersoug, B., Mikkelsen, E. & Osmundsen, T. C. (2021). What's the clue; better planning, new technology or just more money? - the area challenge in norwegian salmon farming. *Ocean Coastal Management*, *199*(1).

Husodo, Z. A. & Vidiapratama, B. (2011). Optimal commodity and cross-currency heding: The case of asean-5-based grain and soft commodity traders.

Jarque, C. M. & Bera, A. K. (1980). Efficient tests for normality, homoscedasticity and serial independence of regression residuals. *Economics Letters*, *6*(3), 255–259.

Kroner, K. F. & Sultan, J. (1993). Time-varying distributions and dynamic hedging with foreign currency futures. *The Journal of Financial and Quantitative Analysis*, *28*(4), 535–551.

Ku, Y.-H. H., Chen, H.-C. & Chen, K.-H. (2007). On the application of the dynamic conditional correlation model in estimating optimal time-varying hedge ratios. *Applied Economics Letters*, *14*(7), 503–509.

Lee, C.-f., Chen, S.-S. & Shrestha, K. (2003). Futures hedge ratios: A review. *The Quarterly Review of Economics and Finance*, *43*(3), 433–465.

Lerøy Seafood Group ASA. (2021). *Annual report 2020*. Retrieved May 1, 2021, from https://www.leroyseafood.com/globalassets/02-documents/rapporter/arsrapporter/arsrapport-2020.pdf

Li, W. K. & Mak, T. K. (1994). On the squared residual autocorrelations in non-linear time series with conditional heteroskedasticity. *Journal of Time Series Analysis*, *15*(6), 627–636.

Lien, D. (2009). The effects of skewness on optimal production and hedging decisions: An application of the skew-normal distribution. *The Journal of Futures Markets*, *30*(3), 278–289.

Liu, Y., Olaussen, J. O. & Skonhoft, A. (2011). Wild and farmed salmon in norway - a review. *Marine Policy*, *35*(3), 413–418.

MOWI. (2020). *Salmon farming industry handbook 2020*. Retrieved February 1, 2021, from https://corpsite.azureedge.net/corpsite/wp-content/uploads/2020/06/Mowi-Salmon-Farming-Industry-Handbook-2020.pdf

MOWI ASA. (2021). *Integrated annual report 2020*. Retrieved May 1, 2021, from https://corpsite.azureedge.net/corpsite/wp-content/uploads/2021/03/Mowi_Integrated_Annual_Report_2020.pdf

Ødegaard, B. A. & Børsum, Ø. G. (2005). Currency hedging in norwegian non-financial firms. *Economic Bulletin*, *2*(05), 133–144.

Oglend, A. (2013). Recent trends in salmon price volatility. *Aquaculture Economics & Management*, *17*(3), 281–299.

Oglend, A. & Sikveland, M. (2008). The behaviour of salmon price volatility. *Marine Resource Economics*, *23*(1), 507–526.

Oglend, A. & Straume, H.-M. (2019). Futures market hedging efficiency in a new futures exchange: Effects of trade partner diversification. *The Journal of Futures Markets*, *40*(4), 617–631.

Palm, F. C. (1996). *Handbook of statistics* (Vol. 14). Elsevier.

Rapp, T. A. & Sharma, S. C. (1999). Exchange rate market efficiency: Across and within countries. *Journal of Economics and Business*, *51*(5), 423–439.

Royal Norwegian Salmon ASA. (2022). *Annual report 2020*. Retrieved May 1, 2021, from https://cdn.sanity.io/files/1gakia31/production/d2f76e0498525c2b0d07831dec33f4d2ed5547ed.pdf

SALMAR ASA. (2021). *Annual report 2020*. Retrieved May 1, 2021, from https://ml-eu.globenewswire.com/Resource/Download/3a869ed4-be25-4ad7-b1c1-39e901bb207f

Samuelson, P. A. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review*, *6*(2), 41–49.

Schwert, G. W. (1989). Tests for unit roots: A monte carlo investigation. *Journal of Business Economic Statistics*, *7*(2), 147–159.

Silvennoinen, A. & Teräsvirta, T. (2009). Modeling multivariate autoregressive conditional heteroskedasticity with the double smooth transition conditional correlation garch model. *Journal of Financial Econometrics*, *7*(3), 373–411.

Tse, Y. K. & Tsui, A. K. C. (2002). A multivariate generalized autoregressive conditional heteroscedasticity model with time-varying correlations. *Journal of Business & Economic Statistics*, *20*(3), 351–362.

Vlaar, P. J. G. & Palm, F. C. (1993). The message in weekly exchange rates in the european monetary system: Mean reversion, conditional heteroscedasticity, and jumps. *Journal of Business Economic Statistics*, *11*(3), 351–360.

Wang, Y., Wu, C. & Yang, L. (2015). Hedging with futures: Does anything beat the naïvehedging strategy? *Management Science*, *61*(12), 2870–2889.

Yun, W.-C. & Kim, H. J. (2010). Hedging strategy for crude oil trading and the factors influencing hedging effectivness. *Energy Policy*, *38*(5), 2404–2408.

Appendix A | Data

A.1 Data characteristics

Table A.1: Out-of-sample test statistics for weekly spot and forward returns.

	JB	ADF	LBQ	LM
Salmon spot	1.66	-2.92***	2.96 (0.937)	7.70***
Salmon forward	254.1***	-4.20***	5.22 (0.734)	28.7***
Currency spot	15.57 ***	-3.86***	3.71 (0.883)	19.74***
Currency forward	6.18 **	-3.67***	7.62 (0.471)	13.15***

Note: Statistics based on significance levels of p*** < 1%, p** < 5% and p* < 10%

A.2 Estimated models

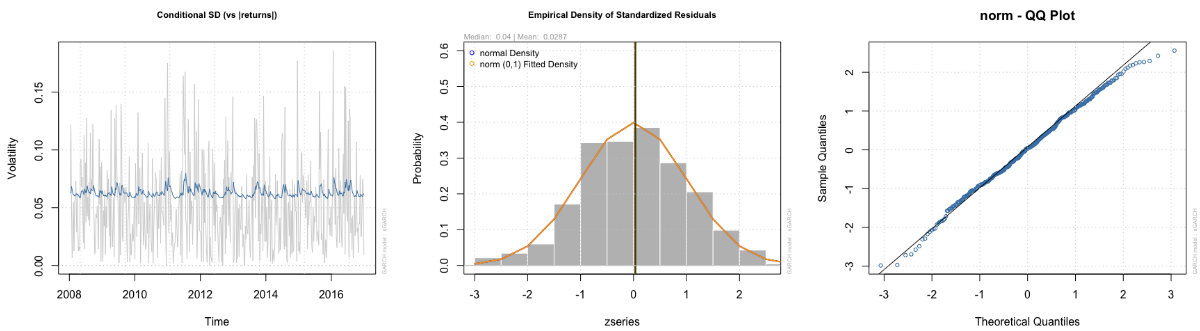

Figure A.1: GARCH(1,1) data plots for salmon spot return series

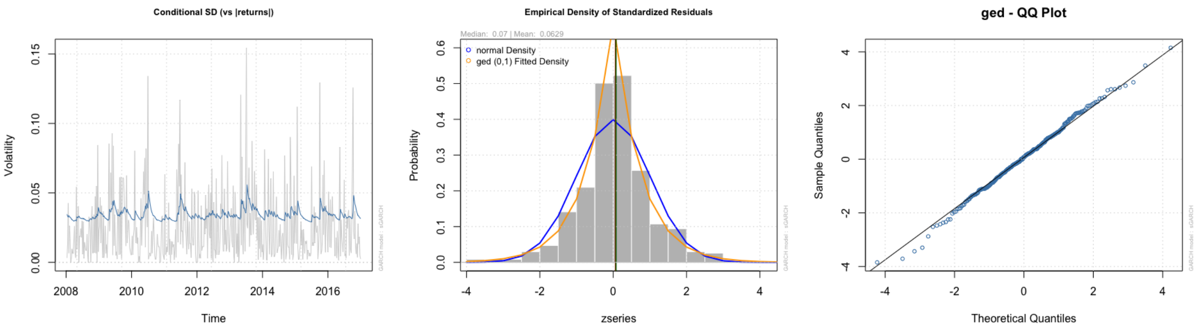

Figure A.2: GARCH(1,1) data plots for salmon 1M forward return series

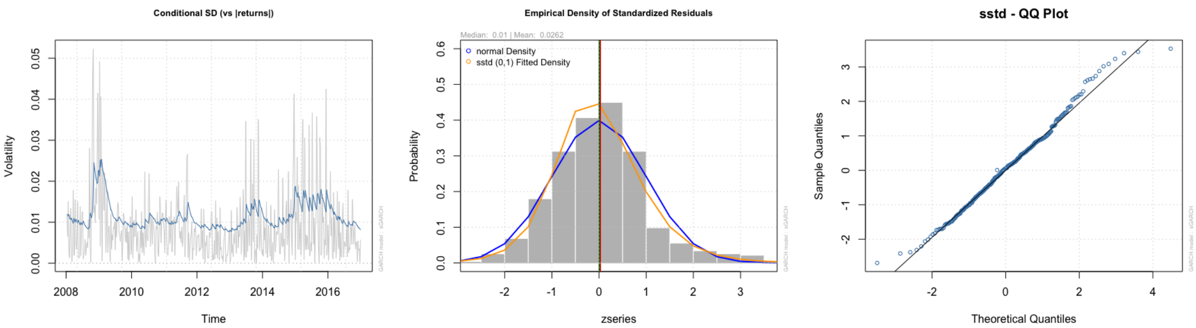

Figure A.3: GARCH(1,1) data plots for currency spot return series

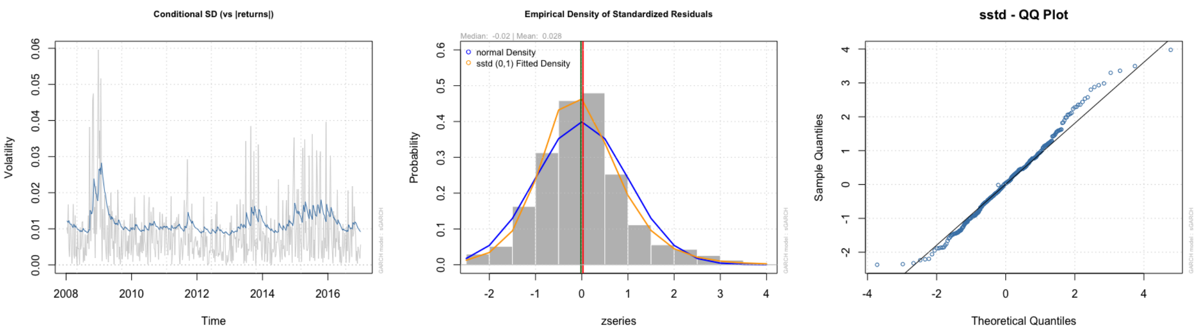

Figure A.4: GARCH(1,1) data plots for currency 1M forward return series

Appendix B | Results

B.1 Hedge ratios

Table B.1: Out-of-sample hedge ratios for dynamic strategies

	Single	Multi
Salmon		
Mean	0.6062	0.5935
SD	0.0861	0.1068
Min	0.2643	0.1674
Max	0.8425	0.8634
Currency		
Mean	0.8557	0.4778
SD	0.0409	0.2426
Min	0.6658	0.0187
Max	0.9651	1.6672

Table B.2: Out-of-sample hedge ratios for threshold strategies.

	Naïve	Dynamic
Salmon		
Mean	0.9737	0.5932
SD	0.1606	0.1271
Min	0.0000	0.0000
Max	1.0000	0.8425

B.2 Return paths

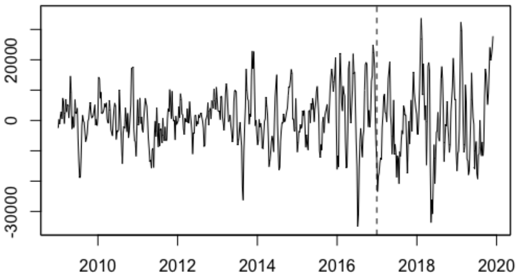

(a) Unhedged

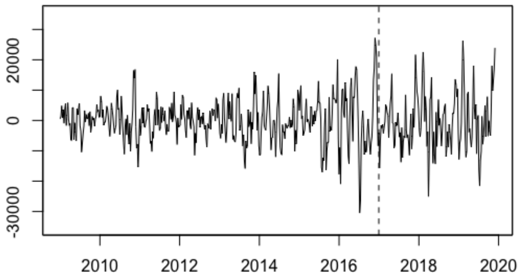

(b) Naïve

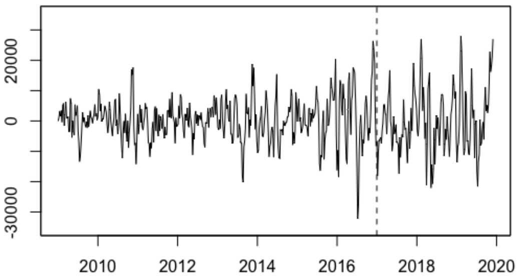

(c) Single hedge

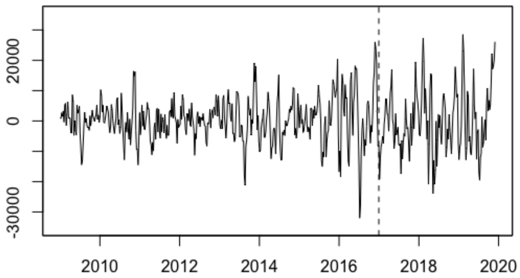

(d) Multi hedge

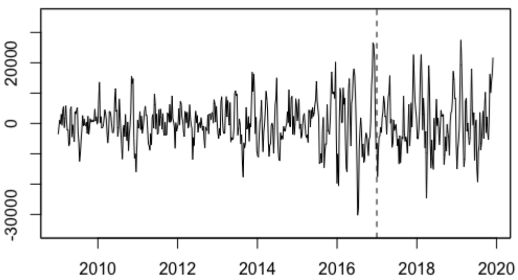

(e) Naïve threshold

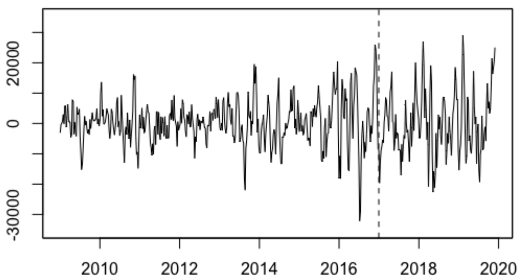

(f) Dynamic threshold

Figure B.1: Return paths for the hedging strategies.
Note: Returns denoted in 1000s NOK

B.3 Robustness tests

Table B.3: Hedge effectiveness for multiple time periods in robustness test.

Hedge effectiveness	Naïve hedge	Single hedge	Multi hedge	Naïve threshold	Dyn. threshold
In-sample					
2009 - 2013	40.58%	38.39%	37.78%	38.75%	35.85%
2010 - 2014	37.90%	33.10%	32.17%	36.87%	30.85%
2011 - 2015	42.80%	37.22%	35.18%	41.97%	33.92%
2012 - 2016	29.99%	26.70%	24.23%	28.77%	22.31%
2013 - 2017	25.92%	22.91%	19.78%	23.50%	19.09%
Mean	35.44%	31.73%	29.83%	33.97%	28.40%
Out-of-sample					
2009 - 2013	43.80%	33.50%	32.55%	43.26%	32.16%
2010 - 2014	45.78%	42.81%	38.16%	37.65%	31.64%
2011 - 2015	-4.40%	7.62%	12.38%	-2.06%	8.23%
2012 - 2016	25.20%	22.35%	20.05%	22.88%	20.34%
2013 - 2017	47.75%	37.07%	22.16%	36.54%	23.18%
Mean	31.63%	29.07%	25.06%	27.65%	23.11%

Note: Returns denoted in 1000s NOK.

Table B.4: Mean return for multiple time periods in robustness test.

Mean return	Unhedged	Naïve hedge	Single hedge	Multi hedge	Naïve threshold	Dyn. threshold
In-sample						
2009 - 2013	201	236	161	146	102	29
2010 - 2014	707	132	206	238	311	348
2011 - 2015	149	-233	-203	-33	76	61
2012 - 2016	1145	-80	367	623	308	693
2013 - 2017	1465	-158	628	851	110	858
Mean	733	-20	232	365	181	398
Out-of-sample						
2009 - 2013	2273	-281	335	527	530	1026
2010 - 2014	-747	-1423	-1399	-1141	-989	-972
2011 - 2015	1039	43	483	590	291	674
2012 - 2016	1968	-686	863	603	-1217	404
2013 - 2017	-3641	-1993	-3215	-2576	-1256	-2550
Mean	178	-868	-587	-400	-528	-284

Note: Returns denoted in 1000s NOK.

NTNU

Norwegian University of
Science and Technology