Master's thesis

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

Aleksander Hyggen Haarstad
Kristian Strypet
Eivind Strøm

Multi-commodity price risk hedging in the Atlantic salmon farming industry

A copula modelling approach

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

NTNU
Norwegian University of
Science and Technology

Preface

This thesis is produced as part of achieving the degree Master of Science at the Norwegian University of Science and Technology, NTNU, in Trondheim, Norway. The field of specialisation is Financial Engineering at Department of Industrial Economics and Technology Management. The thesis is independent work by Aleksander Hyggen Haarstad, Kristian Strypet and Eivind Strøm. The motivation behind the work is based on professional and academic interests.

The Norwegian aquaculture industry has grown fast during the last decades, and global trends suggest increasing demand going forward. Nonetheless, the industry faces growth barriers and volatile profits, and the nature of salmon operations exposes the farmer to multiple risky prices. Today, most salmon farmers hedge less than 30% of their salmon price exposure. Additionally, they do not address important volatility contributors such as fish feed. One reason for this is scarcity of knowledge on how to hedge multiple risks simultaneously, and there is an evident lack of academic research on the subject. We believe more research might have a great impact on the industry, and aim to provide key insights into multi-commodity hedging strategies that salmon producers can employ to reduce their price risk exposure.

We want to profoundly thank our supervisor Associate Professor Maria Lavrutich for deeply valuable discussions, council and review during the semester. Further, we want to thank Assistant Professor at the Norwegian School of Economics, Håkon Otneim, for sharing insights on non- and semi-parametric statistics, and Chief Feed Adviser at Norway Royal Salmon, Kåre Gruven, for sharing insights on the industry. Lastly, we would like to thank our friends and family for their support in this period.

Aleksander Hyggen Haarstad
Kristian Strypet
Eivind Strøm

Trondheim, 27 June, 2020

Sammendrag

Denne studien omhandler styring av prisrisiko i oppdrett av atlantisk laks. I dag sikrer de fleste lakseoppdrettere mindre enn 30% av lakseprisen, mens enkelte ikke sikrer noe. I tillegg er det for øyeblikket ingen oppdrettere som adresserer svingningene i fôrprisene eller råvarene som brukes i produksjonen av fôr. Denne avhandlingen er det første akademiske bidraget til reduksjon av samlet prisrisiko i en lakseoppdrettskontekst. I tillegg til laks tar vi de viktigste råvarene i laksefôret i betraktning; soyamel, hvete og rapsolje.

Vår metode baserer seg på styring av prisrisiko for flere råvarer, kjent som *multi-commodity price hedging*. Antall *futures*-kontrakter som bør kjøpes eller selges per enhet med eksponering i spotmarkedene, kjent som *hedge ratio*, estimeres ved å modellere den flerdimensjonale avhengighetsstrukturen mellom råvareprisene. Dette gjøres med tre forskjellige *copula*-modeller.

Resultatene viser at samlet prisrisiko i lakseoppdrettsnæringen kan reduseres betydelig ved å anvende et fler-råvare-rammeverk med dynamiske *copula*-modeller. Den foreslåtte *rolling window copula multi hedge*-modellen (RWC) reduserer variansen med opptil 53.52%, og utkonkurrerer andre modeller. Dette er modellen som ofrer minst avkastning i forsøket på å redusere prisrisiko. Anvendelsen av flerdimensjonal risikoreduksjon, *multi-commodity hedging*, gir ytterligere risikoreduksjon for kortere perioder, og har en tendens til å forbedre avveiningen mellom risiko og avkastning ved lengre perioder. Videre viser resultatene at å utvide standard flerdimensjonale GARCH-modeller ved anvendelse av *copulaer* reduserer prisrisikoen ytterligere i de fleste tilfeller.

Et annet nøkkelfunn er at periodens lengde har stor innvirkning på hvor mye risikoen kan reduseres. Lakseoppdrettere må foreta en avveining der lengre perioder generelt gir bedre risikoreduksjon og lavere kostnader, men i større grad krever planlegging av fremtidige slaktevolumer. Til slutt foreslår vi et kostnads-effektivitets-mål som understreker viktigheten av å vurdere kostnadene ved risikoreduksjon opp mot hvor mye risikoen reduseres. RWC-modellen er den mest effektive modellen når det kommer til kostnadseffektivitet for lengre perioder. Dette bør være attraktivt for oppdrettsselskapene som for øyeblikket i stor grad foretrekker å være eksponert mot spotprisene i frykt for å gå glipp av positiv avkastning.

Multi-commodity price risk hedging in the Atlantic salmon farming industry: A copula modelling approach

Aleksander Hyggen Haarstad*, Kristian Strypet*, Eivind Strøm*

**Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, 7491, Trondheim, Norway*

27 June, 2020

Abstract

This study addresses the joint input and output price hedging problem for Atlantic salmon farmers. Along with salmon, we consider prices of the three most important commodities in fish feed mixtures; soymeal, wheat, and rapeseed oil. Our approach is based upon multi-commodity price hedging using state-of-the-art of copula models. The results show that joint price risk in the salmon farming industry can be substantially reduced by multi-commodity hedging. The proposed rolling window copula multi-hedge reduces portfolio variance by up to 53.52% and outperforms other models. The use of multi-commodity hedging improves hedging effectiveness for short horizons and tends to improve the risk-return trade-off for longer horizons. Further, our results show that extending the standard multivariate GARCH models by applying copulas increases hedging performance in most cases. Another key finding is that the hedging horizon greatly impacts hedging outcomes. Salmon farmers face a trade-off where longer hedging horizons yield better hedging effectiveness and lower costs but require pre-planned slaughtering volumes to a higher degree. Lastly, we propose a cost-effectiveness measure, highlighting the importance of considering the costliness against the effectiveness of a hedge. By this measure, the RWC model is the most efficient for longer hedging horizons. This is attractive for salmon companies, which currently prefer spot price exposure.

Keywords: Aquaculture, Salmon farming, Salmon feed, Risk management, Multi-commodity hedging, Cross-hedging, Futures, Copulas, GARCH

1 | Introduction

Both the demand and production of Atlantic salmon have been growing fast for the last decades, with Norwegian producers[1] accounting for more than half of world production in 2017 (Brækkan, 2014; Asche et al., 2011; Misund and Asche, 2016; Berge, 2019). At the same time, the last few years have seen a sharp increase in planning and development of land-based salmon production at sites closer to key consumer markets, which could threaten the margins of Atlantic salmon farmers (EY, 2019). Norwegian producers still have a competitive advantage, but the industry faces clear challenges that must be managed to maintain pole position. One of these challenges is the substantial volatility which salmon farming profits feature. The volatility stems from different sources with a significant portion coming from market risk. Most farmers in Norway have acknowledged the importance of managing market risk and try to partially mitigate it by engaging in price risk hedging with exchange traded futures contracts on salmon (Mowi, 2020; SalMar, 2020; Lerøy Seafood Group, 2019; Grieg Seafood, 2019; Norway Royal Salmon, 2020). Such contracts can serve as means for risk transfer from those who wish to reduce risk, typically a salmon farmer, to those with a higher risk appetite.

The salmon price, however, is not the only uncertain factor affecting the profits. Optimising business performance requires successful management of costs and related risks. The main input cost for salmon producers is fish feed (Mowi, 2019). Both the fish feed itself and the commodities in the feed mix feature substantial price volatility, creating an opportunity for the use of novel hedging strategies. Some feed producers have started to offer the feed purchasers to hedge the input commodity prices. However, most farmers seem to be under the perception that, in the long run, costs will outweigh the benefits of hedging exposure in the feed input commodity markets, and thus remain unhedged.[2] Nonetheless, findings in several studies, e.g. Smith and Stulz (1985) and Graham and Smith (1999), suggest that reducing exposure can add significant value. That being the case, there is an evident need for an industry specific examination of joint input and output hedging. This complex hedging problem has received limited attention among practitioners and academics. Potential reasons are a history of satisfactory operating margins, a lack of standardised financial hedging tools such as futures on the feed itself and limited knowledge of the potential and use of financial hedging among industry players.[2]

In this thesis we provide a novel application of multi-commodity hedging where we model the joint risk of input and output price movements. Our first contribution is to provide practical steps towards better risk management practices in the industry by applying advanced techniques to a stylised scenario applicable across different value chain set-ups. We move

[1] The terms salmon producer, salmon farmer and salmon company are used interchangeably in the thesis, referring to the same thing.

[2] This information was revealed in a phone interview with Kåre Gruven, Chief Feed Adviser at Norway Royal Salmon, 19 May 2020.

from traditional output price hedging to joint input and output price hedging. While hedging the output price is already widely examined in the industry, hedging input prices such as feed has been less straightforward in the absence of futures. We show how the contract types used for feed purchases can be exploited to cross-hedge feed price risk. Currently, the full input commodity price risk in salmon feed production is carried by the salmon farmer. This enables the farmer to hedge the feed price risk by taking positions in established exchange traded commodity futures. We provide a first application of copula GARCH[3] models for estimating hedge ratios in the salmon industry. The study examines the share of the salmon production price risk that can be mitigated by simultaneously hedging salmon production input and output price risks. We obtain novel results and find that copula estimation of hedge ratios can significantly improve the risk-return trade-off compared to unhedged portfolios, one-to-one hedged portfolios and portfolios where hedge ratios are estimated by traditional multivariate GARCH (MGARCH) models.

Today, improvement of risk management practices in the salmon farming industry is particularly valuable. Prior to the COVID-19 outbreak, global trends as growing middle class in emerging economies and the industry's relatively low carbon footprint pointed towards strong demand for Atlantic farmed salmon in the years to come (Salmon Facts, 2016; Mowi, 2019). Further, markets have seen sharp drops and increasing volatilities following the COVID-19 outbreak. This has indeed been the case for salmon prices too, dropping close to 30 % between late February and early April 2020, dramatically impacting salmon farming revenues. Increasing volatility in prices for salmon feed input commodities such as soymeal also contributes to higher uncertainty in salmon farming operating margins going forward. This has further exposed the need for better risk management practices in the industry.

Independent from demand trends, the growth of the industry is limited by biological factors (Jensen, 2019). The regulating authorities are concerned about the environmental implications of the industry, such as fish welfare and lice transfer from farmed to wild salmon. To combat this, the government has imposed strict capacity regulations. These limit the growth potential for farmers and the Norwegian salmon industry as a whole (Fiskeridirektoratet, 2020). Under these circumstances, the key to achieving economic sustainability is to ensure profitability in the industry by innovative means that help to tackle existing inefficiencies. Proper management of revenues, costs and associated risks are thus more important than ever before.

The biological nature of the industry leads to periods of higher mortality rates (Hovland, Hopland and Solheimsnes, 2019) and periods of forced excessive slaughtering (Knudsen, 2019). This, together with seasonality in growth and harvesting, results in large variations in salmon supply which feed through into financial markets and contribute to volatile prices (Thyholdt, 2014; Oglend, 2013). The biological factors contribute to profit volatility themselves by affecting the quantity produced. However, management of non-market risks is outside of the scope of this thesis.

Our second contribution is to extend the current literature on hedging salmon farming price risk by applying a multivariate GARCH model to obtain dynamic hedge ratios for both salmon and fish feed commodities. Additionally, we analyse the suitability of GARCH models to capture heteroscedasticity in the time series. Salmon price risk hedging has been subject to extensive academic research, and former studies such as Oglend (2013) has found significant heteroscedasticity in price volatility. Hence, the use of GARCH models is necessary for

[3]Generalised autoregressive conditional heteroskedasticity.

describing volatility and to obtain dynamic hedge ratios. Misund and Asche (2016) examine hedging of salmon spot price exposure by entering salmon futures contracts. They obtain dynamic hedge ratios by applying a bivariate GARCH model, resulting in significant variance reduction. Bloznelis (2018) uses a similar approach, but focuses on relaxing the assumption of known expected prices while at the same time obtaining moderate hedging performance. Our study goes beyond this by examining how to reduce exposure to multiple risks.

The related studies within fish feed hedging are rather limited. Among the few contributions are Vukina and Anderson (1993) and Franken and Parcell (2011). The former studied cross-hedging of fish meal and soybean meal, while the latter provided an extension by considering both soybean meal and corn futures, obtaining improved results. Since the amount of fish meal in modern salmon feed mixes are expected to fall below 10% in the close future, the results have limited value for our study (BioMar Group, 2018). On the other hand, their successful cross-hedging of fish meal suggests similar approaches should be examined for fish feed hedging, which is what we do.

The closest related contribution to hedging input price risk is Haarstad, Strypet and Strøm (2019). Their study is a theoretical contribution to salmon farming input hedging by applying a structural equilibrium model and entering futures contracts on one of the feed input commodities in the absence of feed futures. The hedging effectiveness in terms of lowering the variability in profit was, however, minor. This was explained by the variance in feed prices being dominated by the variance in the salmon prices, which has grown over the past years (Oglend, 2013). We hypothesise that a more successful hedge of input price risk requires a simultaneous hedge of both feed and salmon prices, which we explore in our study.

Our third contribution is to extend the current salmon hedging knowledge base by investigating the potential of state-of-the-art multi-commodity hedging methods. The study is an extension of contributions to output price hedging such as Misund and Asche (2016) and Bloznelis (2018). Multi-commodity hedging has to the best of our knowledge not yet been studied in the context of salmon farming. This thesis fills the gap related to modelling of input hedging in the current aquaculture risk management literature. At the same time, it expands the current knowledge base from solely output price hedging, not only to input price hedging, but further to general price hedging.

We study hedging the joint risk of feed and salmon sales prices within a multi-commodity hedging framework. Even though there is a lack of literature on simultaneous input/output price hedging in an aquaculture business context, similar problems have been examined in other industries. Applications to agriculture are particularly interesting, given the similarities of the two industries. Studies of multi-commodity hedging in cattle farming have yielded good results in terms of reducing profit variability (Anderson et al., 2017) and lowering the risk of big losses (Power et al., 2013).

Power and Vedenov (2009) study the simultaneous hedging of corn (input) and fed cattle (output) for a Texas feedlot operator, which in principle is similar to the hedging problem for a salmon farmer. They show that the hedge ratio for hedging extreme losses is significantly lower than for minimising variance, which is the classical hedging framework. Our study focuses on hedging effectiveness as well. To avoid over-simplifying assumptions of multivariate normality, Power and Vedenov (2009) apply a non-parametric copula (NPC) to model the joint distributions of spot and futures prices for the two commodities considered. One of the main challenges using NPC is the curse of dimensionality where the non-parametric density estimation convergence diminish as dimensions increase (Nagler and Czado, 2016). Given that our practical approach uses multiple commodities, the application of a NPC framework requires new methods and techniques to resolve today's obstacles and challenges. Thus, we use multiple parametric copulas, as applying them for describing the dependence between two variables in many cases can be more effective than linear correlation (Patton, 2006b).

Power et al. (2013) extend the work of Power and Vedenov (2009) by comparing several GARCH techniques in terms of lowering the joint risk of input and output price fluctuations, again for a Texas feedlot operator. They find that the copula GARCH model outperforms both the dynamic conditional correlation (DCC) and Baba-Engle-Kraft-Kroner (BEKK) model in terms of lowering tail risk. Our study explores this in a salmon farming context and further confirm that the results of Power et al. (2013) apply there.

Anderson et al. (2017) also study multi-commodity hedging in the live cattle futures market by comparing hedge ratios of corn under both single- and multi-commodity frameworks. They find that the hedge ratios differ because the multi-commodity hedge ratios of corn are dominated by the cross-dependence between live cattle and corn. This is analogous to the observations of Haarstad, Strypet and Strøm (2019). In their study of the salmon farming industry, the effectiveness of a single-commodity hedge on fish feed are limited as a result of the variance of output prices dominating the variance of feed prices. Anderson et al. (2017) conclude that especially the multi-commodity hedging strategy, as well as the single-commodity hedging strategies, perform better than the non-hedging strategy when considering minimum variance and tail risk criteria. Similarly to Power and Vedenov (2009), Anderson et al. (2017) apply a copula to obtain the joint distribution of spot and futures prices for corn and cattle. Results show that using copula-based methods with GARCH to derive hedge ratios can be more suitable than conventional approaches to computing risk, as these tend to over- or underestimate the risk (Rosenberg and Schuermann, 2006). This suggests that copulas could be useful for modelling hedge ratios in the aquaculture industry, which is confirmed by our study. We show that hedging outcomes are significantly better on most metrics when hedge ratios are estimated by copula methods, compared to when estimated by the DCC model.

From the methodological perspective, our study builds on the seminal contribution in the theory of copulas by Sklar (1959). This study showed that a joint distribution can be transformed into marginal distributions and a copula function which describes the dependence between the variables (Patton, 2006a). Vice versa, marginal distributions can be combined with a copula function to form a joint multivariate distribution, which we utilise in our study. As a measure of dependence between variables, the copula is more informative than linear correlation when the joint distribution of the variables is non-elliptical (Patton, 2006b). The copula approach relaxes the often unrealistic assumption of joint multivariate normality of traditional multivariate GARCH models (Power and Vedenov, 2008; Jondeau and Rockinger, 2006). Copulas can therefore provide realistic joint distributions which can be exploited in risk management by obtaining more realistic GARCH models. Power and Vedenov (2008) describe the extension of copula theory to stochastic processes, i.e. time series, leading to a number of empirical applications of copula theory in financial literature.

The application of copulas are to the best of our knowledge not explored in the aquaculture economics literature. Successful applications in agriculture suggest they have the potential to be useful in an aquaculture economics context too, which we confirm in our study.

The remainder of the paper is structured as follows: Chapter 2 presents the methodology applied in the study. A description of our application to the salmon industry, data and estimated models are presented in Chapter 3. Results are given and discussed in Chapter 4, while Chapter 5 concludes the paper and suggests directions for further research.

2 | Methodology

In this chapter we present the methodological foundation for our models. First, we introduce a set of hedging strategies which create hedging portfolios consisting of simultaneous positions in both spot and futures contracts in several commodity markets. Second, we present four measures to evaluate the effectiveness of the hedged portfolios and capture important differences in performance. Third, we analyse conventional methods used to obtain optimal hedge ratios, being the univariate GARCH(1,1) model and the multivariate dynamic conditional correlation (DCC) model. These will serve as a basis for building more complex models. Fourth, we demonstrate how copulas can be applied as an extension to GARCH models, an approach of increasing popularity in financial econometrics. Lastly, we present three state-of-the-art copula GARCH models, which we use to obtain optimal hedge ratios and hedge the salmon farmer price uncertainty.

2.1 Hedging strategies

Single-hedge

A widely used technique for managing price risk is through hedging with futures contracts. Consider a salmon company with exposure to the price of the commodity produced, and the price of the input commodities required to produce the output. A hedge is then achieved by taking opposite positions in spot and futures markets simultaneously, so that losses resulting from adverse price movements in one market can to some degree be offset by a beneficial movement in the other. The size of the position in futures contracts is determined by the *hedge ratio*, denoted h, which is the number of futures contracts desirable to enter per unit of exposure in the spot market. Following Ederington (1979), risk in this context is measured as the volatility of the company's portfolio of price returns, where the goal is to minimise the portfolio variance by choosing appropriate hedge ratios.

In order to hedge price exposure we consider two commonly employed strategies. The first strategy is the *naïve hedge* where $h = 1$. Implicit in this strategy is a view that the spot and futures market move closely together, and is optimal only if price movements in both markets are proportionate and exactly match each other (Butterworth and Holmes, 2001). However, this is rarely the case. An alternative to the naïve hedge is to find the *optimal hedge ratio*, h^*, which minimises the portfolio variance by taking imperfect correlations into account. The optimal hedge is then estimated under the assumption of constant volatility and correlation, known as *static hedging*.[1] Given that Asche, Misund and Oglend (2016) finds

[1]An estimation of the static hedge ratio is easily undertaken by an OLS-regression of s_t on f_t. Variants of this include rolling-window OLS when extending to dynamic hedge ratios, as employed by Asche, Misund and Oglend (2016).

little difference between the naïve and static optimal hedge, we employ the naïve hedge as our static benchmark. The second strategy and the focus of our thesis is *dynamic hedging* under time-varying volatility and correlation. The goal is then to find the optimal time-varying hedge ratio at time t, conditional on the information set at time $t-1$. Let s_t, f_t denote the spot and futures log price changes (returns), and h_{t-1} the hedge ratio, then the portfolio return r_t is given by Equation 2.1:

$$r_t = s_t - h_{t-1} f_t. \tag{2.1}$$

Following Brooks (2014, p.465-466), we derive the variance minimising dynamic hedge ratio which is given by Equation 2.2:

$$h_t^* = \frac{Cov_t(s_t, f_t)}{Var_t(f_t)}, \tag{2.2}$$

where $Cov_t(s_t, f_t)$ is the conditional covariance between spot and futures returns at time t and $Var_t(f_t)$ is the conditional variance of the futures returns at time t. The problem of finding the optimal time-varying hedge ratios then becomes estimation of the conditional-variances and covariances for spot and futures price returns in the portfolio.

Multi-hedge

While the optimal hedge ratio of Equation 2.2 holds when considering the spot and futures price returns of a single commodity, it is not necessarily optimal when considering a multi-commodity problem with both input and output. Using a similar approach to Anderson et al. (2017), we tackle this by defining a *single-hedge* where commodities are considered separately, and a *multi-hedge* which exploits the dependency between the different commodities.

First, consider the case where the return on the company's portfolio of commodities is a combination of the variance minimising portfolios of each commodity, hedged independently with hedge ratios as given by Equation 2.2. Here, the hedger assumes that when each commodity is hedged separately, the combination results in a portfolio that reduces overall risk. In this setting, the dependency between different commodities is not considered and there are no opportunities for cross-hedging (Anderson and Danthine, 1981). This necessarily prevents speculative positions when the spot and futures markets are positively correlated.[2] We denote the vector of optimal dynamic hedge ratios when considering i commodities hedged separately as $\mathbf{h}_{S,t} = \{h_{1,t}, ..., h_{i,t}\}$, which we refer to as the single-hedge ratio.

Second, we consider the combined returns on the portfolio of all commodities in a multi-commodity setting, following the hedging framework of Anderson and Danthine (1981). In this framework, the commodities are considered in unison. This implies that unfavourable movements in one commodity spot price can be more effectively offset by movements in a different commodity price rather than just the corresponding commodity futures price. This entails both cross-hedging and speculative positions in different markets to obtain the combined minimum variance portfolio. Furthermore, it depends on the spot commodity quantities, implying that exposures are weighted higher. We denote the vector of optimal hedge ratios $\mathbf{h}_{M,t} = \{h_{1,t}, ..., h_{i,t}\}$ and will henceforth refer to it as the multi-hedge ratio, given

[2] A speculative position entails going long (or short) both the corresponding spot and futures market simultaneously, effectively *increasing* the exposure.

by Equation 2.3:[3]

$$\mathbf{h}^*_{\mathbf{M,t}} = \left[diag(\mathbf{Q})\right]^{-1}\sum\nolimits_{FF}^{-1}(t)\sum\nolimits_{FP}(t)\mathbf{Q}, \tag{2.3}$$

where $\sum_{FF}(t)$ is the $(m \times m)$ variance-covariance matrix of futures prices, $\sum_{FP}(t)$ is the $(m \times m)$ variance-covariance matrix of spot and futures prices, $\mathbf{Q}$ is a $(m \times 1)$ vector of the quantities of spot commodities[4] and $diag(\mathbf{Q})$ is a diagonal matrix with $\mathbf{Q}$ on the main diagonal (Fackler and McNew, 1993).

To measure the effectiveness of the different hedging strategies, we apply four measures defined in the following section.

2.2 Measures of hedging efficiency

We consider two main aspects when comparing the effects of different hedging strategies: return and risk, each with two accompanying measures.

The effects of hedging return are measured in two ways. First is mean return which we estimate from a portfolio with historical average returns (French and Fama, 1989; Fama, 1990; Fama and French, 1992). Mean return is the profit or loss the company historically would have received by applying the respective hedging strategies. Second, given that different hedging strategies involve different sized positions in the futures market, we compare the cost of the hedges by computing the transaction costs associated with each hedge. Transaction costs play an important role regarding choosing the optimal hedging strategy. Less frequent rebalancing is cheaper, yet more risky, whereas frequent rebalancing is more expensive, but less risky (Toft, 1996).

The hedging effect on risk is measured by hedge effectiveness (HE) and expected shortfall (ES). When the goal is to minimise the variance of returns, HE is measured as the percentage reduction of variance in the hedged portfolio against the unhedged portfolio, given by Equation 2.4 (Ederington, 1979):

$$\text{Hedge effectiveness} = 1 - \frac{Var(\text{Hedged portfolio})}{Var(\text{Unhedged portfolio})}. \tag{2.4}$$

Tail risk refers to the most extreme downside losses, of magnitude to potentially do great damage in an economic perspective. As a proxy measure for tail risk and financial distress, we employ ES. ES measures the average loss in the worst $\alpha = A\%$ cases, given by Equation 2.5. While value-at-risk (VaR) is often employed for this purpose, it is not sub-additive, nor does it consider the severity of losses in worst case scenarios. ES is therefore used as a more coherent measure of tail risk (Acerbi and Tasche, 2001).

$$ES^{\alpha}(X) = \left(-\frac{1}{\alpha}\right)\left(\mathbb{E}\left[X\mathbb{1}_{X\le x^{\alpha}}\right] - x^{\alpha}\left(\mathbb{P}[X \le x^{\alpha}] - \alpha\right)\right). \tag{2.5}$$

ES can be simplified to tail conditional expectation (TCE) when the probability distributions are continuous:

$$TCE^{\alpha}(X) = -\mathbb{E}\left\{X|X \le x^{\alpha}\right\}. \tag{2.6}$$

[3]The original framework formalised by Fackler and McNew (1993) has been extended from the static to the time-varying case by applying t subscripts.

[4]Positive (negative) quantities correspond to long (or short) positions.

2.3 GARCH models

In order to obtain time-varying hedge ratios and capture important characteristics such as heteroscedasticity, dependence between variables and tail behaviour in our financial time series data, we estimate GARCH models. In what follows we present the standard GARCH(1,1) model and the DCC model, which serve as a basis for the more complex models introduced later.

GARCH(1,1) model

Consider a time series of commodity prices with a sample of T observations. Letting P_t define the time series evaluated at time t, the continuously compounded return r_t is defined as the log-change, given by Equation 2.7:

$$r_t = lnP_t - lnP_{t-1}. \tag{2.7}$$

We let the unconditional mean and variance be denoted by μ and σ^2. Then, the conditional mean and variance, μ_t and σ_t^2, can be written as[5]

$$\begin{aligned} \mu_t &= \mathbf{E}[r_t|\Im_{t-1}], \\ \sigma_t^2 &= \mathbf{E}[(r_t-\mu_t)^2|\Im_{t-1}], \end{aligned} \tag{2.8}$$

where $\Im_{t-1}$ denotes the information available at $t-1$. The return at time t is then given by

$$r_t = \mu_t + \sigma_t\epsilon_t, \quad \text{where} \quad \epsilon_t \sim g(0,1,\theta), \tag{2.9}$$

ϵ_t is the standardised residual at time t and $g(0,1,\theta)$ is the assumed conditional distribution with distributional parameters θ. While the original GARCH model assumes ϵ_t to be standard normal, we also consider the generalised error distribution (GED), the Student's t distribution and the skewed t distribution of Fernández and Steel (1998).

The conditional variance is modelled by

$$\sigma_t^2 = \omega + \beta h_{t-1}^2 + \alpha u_{t-1}^2, \tag{2.10}$$

where $u_t = \sigma_t\epsilon_t$ and ω, β, α are the parameters of the process. With the above specification, the unconditional variance of u_t is given as $\text{Var}(u_t) = \frac{\omega}{1-(\alpha+\beta)}$. To ensure stationarity, we require the restriction $(\alpha+\beta) < 1$ (Brooks, 2014, p. 430). While the GARCH(1,1) model can be extended to a GARCH(p,q) model, the (1,1)-specification is generally sufficient and has been found to perform well compared to higher order models when an appropriate distribution for ϵ_t is specified (Brooks, 2014; Hansen and Lunde, 2005).

Next we present the DCC model, which will be the baseline multivariate GARCH model for modelling conditional covariances.

[5]While conditional variance is commonly denoted as h_t in financial literature, we use the σ_t^2 notation to avoid confusion with the hedge ratio h.

DCC model

In order to capture time-varying correlations between time series used in dynamic hedging, a multivariate GARCH (MGARCH) model is needed. There exist many different specifications of MGARCH in the literature (Bauwens, Laurent and Rombouts, 2006). One of the main challenges with MGARCH models however, is the curse of dimensionality, where the number of parameters increase rapidly with the dimensionality of the models (Caporin and McAleer, 2014). This is a relevant challenge for commodity processors such as a salmon farmer, hedging against the risk of multiple commodities. For example, the general BEKK model in the case of four commodities (8 time series), would require jointly estimating 164 parameters.[6] Therefore, we will focus on specifications that allow us to use higher dimensions such as the widely used DCC model proposed by Engle (2002) and Tse and Tsui (2002).

For this model family the conditional covariance estimation is simplified by estimating GARCH(1,1) models for each commodity. The transformed residuals from each commodity is used to estimate a conditional correlation estimator which is then used to modify the standard errors for the correlation parameters. The variance-covariance matrix H_t is defined as

$$\mathbf{H}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t, \tag{2.11}$$

where D_t is a diagonal matrix containing the conditional standard deviations obtained from Equation 2.10 for each individual series, and R_t is the conditional correlation matrix. Both D_t and R_t vary over time, producing a new variance-covariance matrix for each time step, differentiating DCC from the constant conditional correlation (CCC) model.[7]

DCC models do not come without shortcomings. One of them being the requirement in the maximum likelihood estimation (MLE) procedure that the standardised residuals follow the multivariate normal distribution. This may not be consistent with financial data, which often can contain features such as skewness and excess kurtosis. To address this issue, we apply a copula approach (Patton, 2006b).

2.4 Copulas

An n-dimensional copula, C, is a distribution function with uniformly distributed margins in $[0,1]$. Sklar (1959) showed that any joint distribution function F of the random vector $X = (x_1, ..., x_n)$ with margins $G_1(x_1), ..., G_n(x_n)$, can be decomposed as follows:

$$F(x_1, ..., x_n) = C\big(G_1(x_1), ..., G_n(x_n)\big), \tag{2.12}$$

with the copula C being uniquely determined in $[0,1]^n$ and assuming F with sufficiently smooth margins for derivatives to exist, obtained as

$$C(x_1, ..., x_n) = F\big(G_1^{-1}(u_1), ..., G_n^{-1}(u_n)\big). \tag{2.13}$$

[6] Using the R package mgarchBEKK by Schmidbauer, Roesch and Tunalioglu (2016).

[7] For more details on how the DCC model is estimated with the maximum likelihood estimator (MLE) method, see Engle (2002) and Tse and Tsui (2002).

Conversely, the joint density function f, and copula density c, is obtained as

$$\begin{aligned} f(x_1,...,x_n) =& c\big(G_1(x_1),...,G_n(x_n)\big) \times \prod_{i=1}^{n} g_i(x_i), \\ \text{where} \quad c(u_1,...,u_n) =& \frac{\partial^n C\big(u_1,...,u_n\big)}{\partial u_1 \cdot ... \cdot \partial u_n}, \end{aligned} \tag{2.14}$$

and $u_i = G_i(x_i)$ are the uniform observations in $[0,1]$ transformed by the probability integral transform (PIT), using its series marginal distribution. The resulting u_i are typically referred to as pseudo-observations. An alternative to using the marginal distribution is using the empirical distribution function (EDF) given by Equation 2.15, to obtain pseudo-observations.[8] Estimation of c using the EDF is shown to be consistent, asymptotically normal and fully efficient under the assumption that X is i.i.d. (Genest, Ghoudi and Rivest, 1995).

$$G_i(x) \equiv \frac{1}{T+1} \sum_{t=1}^{T} \mathbf{1}\big\{\hat{x}_{i,t} \leq x\big\}. \tag{2.15}$$

Different parametric copula functions are often referred to as *copula families*. The copula families used in this thesis belong to two main categories: (1) elliptical copulas with symmetric dependency structures and (2) Archimedian copulas with asymmetric dependency structures. We use the normal and Student's t copula of the elliptical category and the Clayton, Gumbel, Frank and Joe copula from the Archimedian category. In the following, we present the density functions and discuss properties of the normal, Student's t and Gumbel copula which are commonly applied to financial data.[9] Figure 2.1 displays normalised contour plots for the normal, Student's t, Gumbel and survival Gumbel copula with parameters given in the upper row. The bottom row displays $N = 300$ simulated pseudo-observations from the respective copulas.

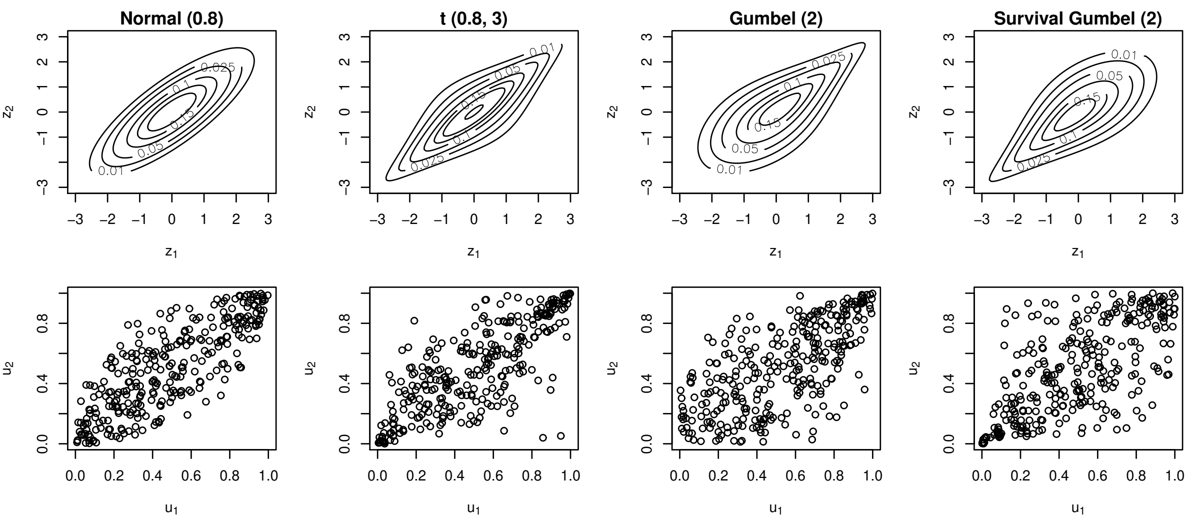

Figure 2.1: Contour plots of normal, Student's t, Gumbel and survival Gumbel copula.

[8]The uniform pseudo-observations can in theory be obtained by the PIT using respective marginal distributions. In our experience, the EDF often produces more uniform margins resulting in more accurate copula estimates unless the marginal distributions are perfectly specified.

[9]Similar details on the Clayton, Frank and Joe copula are attached in Appendix A.2.

The normal copula, commonly referred to as the *Gaussian* copula, is constructed from the multivariate standard normal distribution. The bivariate functional form of the normal copula density is given as

$$c_N(u_i,u_j;\rho) = \frac{1}{\sqrt{1-\rho^2}}\exp\left\{-\frac{\rho^2\left(\Phi^{-1}(u_i)^2+\Phi^{-1}(u_j)^2\right)-2\rho\Phi^{-1}(u_i)\Phi^{-1}(u_j)}{2(1-\phi^2)}\right\}, \tag{2.16}$$

for $\rho \in (-1,1)$ where Φ^{-1} is the inverse cumulative distribution function (CDF) of a standard normal random variable.

Similar to the normal copula, the Student's t copula is based on the standard t distribution. Hence, the t copula generalises normal copula by allowing non-zero dependence in the extreme tails, as seen in Figure 2.1.[10] The bivariate Student's t copula density is given by Equation 2.17:

$$c_T(u_i,u_j;\rho,\nu) = \frac{\Gamma((\nu+2)/2)/\Gamma(\nu/2)}{\nu\pi\, t(x;\nu)\, t(y;\nu)\sqrt{1-\rho^2}}\left(1+\left(\frac{x^2+y^2-2\rho xy}{\nu(1-\rho^2)}\right)^{-\frac{\nu+2}{2}}\right), \tag{2.17}$$

for $\rho \in (-1,1)$ and $\nu > 0$, where $t(\cdot;\nu)$ is the probability distribution function (PDF) of a Student's t random variable with ν degrees of freedom, $x \equiv T^{-1}(u_i;\nu)$, $y \equiv T^{-1}(u_j;\nu)$ and $T^{-1}(\cdot;\nu)$ is the inverse CDF of a Student's t random variable with ν degrees of freedom.

Lastly, the Gumbel copula is an Archimedian copula and therefore allows for asymmetrical tail dependence. Its bivariate density is given by Equation 2.18:

$$\begin{aligned} c_G(u_i,u_j;\rho) =& C_G(u_i,u_j;\rho)(u_iu_j)^{-1}\frac{(\tilde{u}_i\tilde{u}_j)^{\rho-1}}{(\tilde{u}_i{}^{-\rho}+\tilde{u}_j{}^{-\rho})^{2-\frac{1}{\rho}}}\left(\left(\tilde{u}_i{}^{-\rho}+\tilde{u}_j{}^{-\rho}\right)^{\frac{1}{\rho}}+\rho-1\right), \\ \text{where}\quad & C_G(u_i,u_j;\rho) = e^{-(u_i^{\rho}+u_j^{\rho})^{\frac{1}{\rho}}}, \end{aligned} \tag{2.18}$$

for $\rho \in [1,\infty)$. Specifically, the Gumbel copula displays greater dependence in the positive tail than in the negative. Furthermore, Archimedian copulas can be be rotated by 90 degree quadrants to exhibit different tail dependencies. For example, a 180 degree rotation, which is denoted as a *survival* copula,[11] would produce the mirrored asymmetry. Hence, the survival Gumbel displays greater dependence in the negative tail than the positive, which is more appropriate for financial data and depicted in Figure 2.1. 90 and 270 degree rotations are necessarily required to capture negative dependence when considering Archimedian copulas (Brechmann and Schepsmeier, 2013).[12]

Conditional copulas

The copula theory by Sklar (1959) was developed for applications where data is assumed to be i.d.d., and hence not typical time series data. Patton (2006a) proves that copulas can be applied to the case of serially-dependant data if the latter satisfy the Markov Property; Stating that future realisations of a stochastic process conditioned on both past and present

[10]When $\nu \to \infty$, Student's t copula approaches the normal copula.

[11]A definition of survival copulas and joint survival functions can be found at Cherubini (2004, p.75-80).

[12]Details on the 90, 180 and 270 degree rotations for the Archimedian copulas are presented in Appendix A.2.

states depends only on the present state and not on the entire information set $\Im_{t-1}$. I.e. $\mathbf{P}(X_t = x_t|\Im_{t-1}) = \mathbf{P}(X_t = x_t|X_{t-1} = x_{t-1})$. Although the Markov Property does not hold for typical financial time series, it is satisfied by the innovations of fitting a GARCH(1,1) model, assuming the conditional distribution is correctly specified. Therefore, the bivariate conditional joint density of standardised GARCH residuals $\epsilon_{i,t}$, $\epsilon_{j,t}$ at time t can be written as

$$f_t(\epsilon_{i,t},\epsilon_{j,t}|\Im_{t-1}) = c_t\Big(G_{i,t}(\epsilon_{i,t}|\Im_{t-1}), G_{j,t}(\epsilon_{j,t}|\Im_{t-1})\Big|\Im_{t-1}\Big) \times g_{i,t}(\epsilon_{i,t}|\Im_{t-1}) \times g_{j,t}(\epsilon_{j,t}|\Im_{t-1}),$$
$$\text{where} \quad c_t(u_{i,t}, u_{j,t}|\Im_{t-1}) = \frac{\partial^2 C_t\big(u_{i,t}, u_{j,t}|\Im_{t-1}\big)}{\partial u_{i,t}\partial u_{j,t}}, \tag{2.19}$$

$g_{i,t}(\epsilon_{i,t}|\Im_{t-1})$ is the conditional marginal density of $\epsilon_{i,t}$ and $g_{j,t}(\epsilon_{j,t}|\Im_{t-1})$ the conditional density of $\epsilon_{j,t}$.

Parameters of Equation 2.19 are estimated by a two-stage maximum likelihood framework in which parameters for the density functions of $\epsilon_{i,t}$, $\epsilon_{j,t}$ and parameters for the copula function are estimated in two steps.[13] In the first stage the marginal densities are estimated by the fitting of a GARCH(1,1) model for each random variable. In the second stage the parameters of the copula is estimated by maximising the log-likelihood function given the estimates of stage one. Estimates have been shown to be consistent and asymptotically normal under standard conditions (Patton, 2006a).

After estimating parameters for the conditional copula, the copula is combined with the conditional marginal distributions to obtain the conditional joint density from which the conditional covariance can be generated by numerical integration using Equation 2.20:[14]

$$\sigma^2_{ij,t} = \sigma_{ii,t}\sigma_{jj,t}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\epsilon_{i,t}\epsilon_{j,t}f(\epsilon_{i,t}\epsilon_{j,t}|\Im_{t-1})d\epsilon_{i,t}d\epsilon_{j,t}, \tag{2.20}$$

where $\sigma^2_{ij,t}$ is the conditional covariance, $\sigma_{ii,t}$, $\sigma_{jj,t}$ the conditional standard deviations obtained from Equation 2.10 and $f(\epsilon_{i,t}\epsilon_{j,t}|\Im_{t-1})$ is the conditional joint distribution of standardised residuals obtained from Equation 2.19. The resulting conditional variances and covariances are used to compute the optimal hedge ratios of Equations 2.2 and 2.3.

In the next section we present the copula models we use to obtain the conditional joint density of Equation 2.20.

2.5 Copula GARCH models

In the following section we present three copula GARCH models which incorporate the copula theory introduced in the last section. As opposed to GARCH models, copula based methods allow for more flexible modelling of the dependence structure between variables. Most notably, one of the main strengths is the specification of the multivariate distribution by considering the marginal distribution and dependence structure separately.

[13] Details on the log-likelihood function and estimators can be found at Patton (2020a) and Patton (2020b).

[14] In practice, the integral is estimated by simulation from the conditional joint density, as the numerical integration accuracy has proven to be problematic in some cases.

Copula-DCC model

The first copula model we consider is an extension of the standard DCC model. Namely the copula-DCC model (C-DCC), as described in Ghalanos (2019b). Given that the DCC model implicitly assumes a normal copula by assuming a multivariate normal distribution, a relatively simple extension can be made to change to the Student's t copula and make it time varying. In effect, by transformation using Sklar's theorem, we relax the assumption that the distribution of the conditional marginals are standard normal and allow a non-normal dependency structure.

Assume we have $i = 1, ..., n$ conditional marginal distributions estimated in the first stage by GARCH(1,1) processes, where G_i is the conditional CDF of the i^{th} margin. Furthermore, the dependence structure of the margins is assumed to follow a Student's t copula with conditional correlation $\mathbf{R}_t$ and constant shape parameter η. $\mathbf{R}_t$ is assumed to follow a DCC model as described previously. The conditional density at time t is then given by

$$c_t(u_{1,t}, ..., u_{n,t}|\mathbf{R}_t, \eta) = \frac{f_t\Big(G_1^{-1}(u_{1,t}|\eta), ..., G_n^{-1}(u_{n,t}|\eta)\Big|\mathbf{R}_t, \eta\Big)}{\prod_{i=1}^{n} f_i\Big(G_i^{-1}(u_{i,t}|\eta)|\eta\Big)}, \tag{2.21}$$

where $u_{it} = G_{i,t}(x)$ is the PIT of each series by its EDF,[15] $G_i^{-1}(u_{i,t}|\eta)$ is the quantile transformation of the pseudo-observations given the common shape parameter, $f_t(\cdot|\mathbf{R}_t, \eta)$ is the multivariate density of the Student's t distribution and $f_i(\cdot|\eta)$ is the univariate margins of the multivariate t distribution with common shape parameter η.

As a result, the joint density of the two-stage estimation is given by

$$f_t(\mathbf{r}_t|\mathbf{h}_t, \mathbf{R}_t, \eta) = c_t(u_{i,t}, ..., u_{n,t}|\mathbf{R}_t, \eta) \prod_{i=1}^{n} \frac{1}{\sigma_{i,t}} g_{i,t}(\epsilon_{i,t}|\theta_i), \tag{2.22}$$

where $\epsilon_{i,t} \sim g_i(0, 1, \theta_i)$ are the standardised residuals of the stage one estimation with appropriate conditional distributions and parameters θ_i. Conditional covariances are obtained from the conditional joint density by simulation (Ghalanos, 2019b).

Time-Varying Copula model

The second model is a time-varying copula model (TVC), initially proposed in Patton (2006a) and Patton (2006b). In this model, the time variation in the conditional copula parameter is elected to follow a GARCH-like process in which the correlation parameter at time t is the function of a constant ω, the lagged correlation β, and some forcing variable α. Following Patton (2006a), the time-varying parameter for the normal, Student's t and Gumbel copula are modelled as:

$$\begin{aligned}
\text{Normal:} \quad & \rho_t = \Lambda\left(\omega_N + \beta_N \rho_{t-1} + \alpha_N \frac{1}{n} \sum_{k=1}^{n} \Phi^{-1}(u_{i,t-k})\Phi^{-1}(u_{j,t-k})\right), \\
\text{Student's t:} \quad & \rho_t = \Lambda\left(\omega_T + \beta_T \rho_{t-1} + \alpha_T \frac{1}{n} \sum_{k=1}^{n} T^{-1}(u_{i,t-k}; \nu) T^{-1}(u_{j,t-k}; \nu)\right), \\
\text{Gumbel:} \quad & \theta_t = \kappa\left(\omega_G + \beta_G \theta_{t-1} + \alpha_G \frac{1}{n} \sum_{k=1}^{n} |u_{i,t-k} - u_{j,t-k}|\right),
\end{aligned} \tag{2.23}$$

[15]Again, it is also possible to use the conditional marginal distribution for the probability integral transform.

where Φ^{-1} is the inverse CDF of a standard normal random variable and $T^{-1}(\cdot;\nu)$ is the inverse CDF of a Student's t random variable. We use a logistic transformation to ensure $\rho_t \in [-1,1]$ with $\Lambda(x) = \frac{1-e^{-x}}{1+e^{-x}} = tanh(\frac{x}{2})$. The function $\kappa(x) = 1 + x^2$ is used to ensure $\theta_t \in [1,\infty)$. The Equations of 2.23 are estimated by maximum likelihood.[16]

The model therefore consists of estimating the Equations of 2.23 for each time series pair i, j, and selecting the best fitting model by AIC. The resulting conditional copula c_t with time-varying parameter ρ_t or θ_t is combined with the conditional marginals by Equation 2.19. The conditional covariances $\sigma^2_{ij,t}$ are then estimated by Equation 2.20.

Rolling Window Copula model

Lastly, we propose a rolling window copula model (RWC) which allows for time variation in both the copula dependence parameter and the parametric copula family. This contrasts with the DCC, C-DCC and the TVC model which assumes that the copula family modelling the distribution is constant over the sample.[17] Additionally, the DCC and C-DCC models specifically have the disadvantage that they in the multi-dimensional setting assume the same parametric copula for any pair of random variables in the model. It seems questionable to assume that all variable pairs can be modelled appropriately by the same elliptical copula, such as the normal or Student's t assumed in this case. To relax these assumptions, we consider each variable pair individually and select the best fitting conditional bivariate copula c_t among $C_{Families}$ = {normal, Student's t, Clayton, Gumbel, Frank, Joe}.[18] Therefore, we allow different dependency structures between different variable pairs that additionally vary through time.

More specifically, the model estimation is done in two stages. The first stage consists of estimating the conditional marginal distributions by GARCH(1,1) models with appropriate distributions as described in Section 2.3. The second stage consists of estimating conditional copulas c_t between variable pairs, which ultimately are used to obtain the conditional covariances of Equation 2.2 and 2.3. We apply the estimation procedure to a moving window of N observations. I.e. $t = \{1,\ldots,N\}$ observations are used to estimate densities at $t = N$, $t = \{2,\ldots,N+1\}$ for densities at $t = N+1$, and so forth. This allows both the conditional copula family and the dependency parameter to be time-varying. For each conditional covariance $\sigma^2_{ij,t}$, we estimate the bivariate parametric copula of each family in C_{Families} given $u_{i,t}$, $u_{j,t}$, by maximum likelihood. Pseudo-observations $u_{i,t}$, $u_{j,t}$ are obtained by the EDF of Equation 2.15 using the standardised residuals $\epsilon_{i,t}, \epsilon_{j,t}$ estimated in the first stage. The best fitting copula family is then selected by the AIC criterion, given by

$$\text{AIC} = -2\sum_{t=1}^{n} ln\big[c(u_{i,t}, u_{j,t}|\theta_c)\big] + 2m, \tag{2.24}$$

where $m = 1$ for one parameter copulas and $m = 2$ for two parameter copulas, i.e. the Student's t copula. Continuing, we combine the selected copula with the respective conditional marginal distributions $G_{i,t}(\epsilon_{i,t}|\Im_{t-1})$ and $G_{j,t}(\epsilon_{j,t}|\Im_{t-1})$ to obtain the conditional joint density, i.e. Equation 2.19. The conditional covariance is obtained by Equation 2.20, given the conditional joint distribution and the conditional variances.

[16]See Patton (2006a) and Patton (2006b) for more details.

[17]The DCC model implicitly assumes a normal copula when using the multivariate normal distribution.

[18]Including 90, 180 and 270 degree rotations of the Clayton, Gumbel and Joe copula.

3 | Estimation

In the following chapter we describe a stylised problem for a salmon producer. We present the hedging context followed by a description of the data. Lastly, we show the estimated models.

3.1 Application to salmon farming

In this section, we construct a conceptual hedging framework tailored to the salmon farming industry, considering both input and output commodity prices. We consider a hypothetical well-established salmon farming company which has an objective of reducing the price exposures of its operations.[1] Specifically, we make the following assumptions when constructing the salmon production price hedge:

First, the company has a harvest quantity equal to the average of the Norwegian companies listed in the OSLO Seafood Index (OSLSFX). This is equal to 160 000 tons according to the companies 2018 annual reports, and is comparable to companies such as Lerøy Seafood Group and SalMar. The hypothetical company may have several production sites, but operates solely in Norway.

Second, we assume the company has biomass assets in all stages of the salmon production cycle, and harvests salmon continually. The same price is realised for all salmon sold within the same week. The average weekly quantity sold is 3 000 tons.

Third, in practice, the salmon farmer has some flexibility in choosing when to harvest,[2] but for the purpose of our study we assume a constant volume is slaughtered and sold every week. This is consistent with related studies such as Anderson et al. (2017).

Fourth, in practice, feed consumption varies with sea temperatures, but for the purpose of our study we assume a constant consumption throughout the year. The fact that salmon demands more nutrition as it grows does not need to be taken into account as the company's production sites contain salmon in all stages of the production cycle. A feed conversion ratio of 1.1 (Mowi, 2019) implies that the company uses 3 300 tons of feed per week. Feed is bought weekly[3]and we assume that price movements in the feed input commodity markets affect feed prices the same week.

Further, the company aims to reduce the exposure to risk associated with prices of both fish feed and salmon. As means for risk mitigation such as futures contracts on the fish feed

[1]The production cycle for Atlantic salmon lasts for roughly three years (Seafish, 2012).

[2]Incentives to rush (delay) harvest can be to exploit (wait for) favourable prices or periods of year known for faster biomass growth.

[3]This information was revealed in a phone interview with Kåre Gruven, Chief Feed Adviser at Norway Royal Salmon, 19 May 2020.

itself are unavailable, the company proceeds by cross-hedging individual fish feed components. The assumption that fish feed can be cross-hedged by its components is reasonable, as the feed producers traditionally have sold feed on contracts that transfer the full risk exposure to the purchasing party, such as cost-plus type contracts (Mowi, 2019; BioMar Group, 2019). The worlds largest salmon production company, Mowi, controls fish feed costs by upstream integration of the fish feed production. In this case the company can better manage production risks and remove the profit margin of the feed producer. However, they are still exposed to the underlying input commodity price, making our hedging framework applicable across multiple value chain set-ups.

According to Aas, Ytrestøyl and Åsgård (2019), the main components of fish feed are soy protein concentrate (19.0%), wheat and wheat gluten (17.9%), rapeseed and camelina oil (19.8%) and fish meal and oil (24.9%). Other components individually account for less than 4% of the total feed mixture and are hence considered negligible as contributors to feed price volatility. As stated, fish meal and fish oil are important ingredients. However, there are no futures contracts available on these commodities. A hedge of price risk could therefore only be achieved by cross-hedging with e.g. soymeal, similar to Vukina and Anderson (1993). Even though this is possible from a theoretical perspective, it is difficult to do in practice, especially since the correlation between fish meal and soymeal has been lower during recent years (Franken and Parcell, 2011). Given the decreasing share of fish meal and fish oil in modern feed compositions (BioMar Group, 2018), we choose to leave this out in the analysis. We assume the price exposure for each unit of fish feed to be equivalent to 20% exposure to prices of soymeal, wheat and rapeseed oil respectively, as feed compositions to some degree can vary (Mowi, 2019). A feed conversion ratio of 1.1, implies that 0.22 kg of soymeal, wheat and rapeseed oil (0.66 kg total) are required for each 1 kg of salmon produced. Further, it is important to consider that agricultural futures contracts are different from salmon futures as they have fixed sizes. E.g. soymeal and wheat have full-contracts of 100.0 tons and 5 000.0 bu[4] respectively (Parcell and Franken, 2011). For the purpose of our analysis, we assume the contract sizes are sufficient for the company's hedging requirements.

In order to reduce spot price risk, it is desirable to enter a futures contract which price changes are highly correlated to the spot price changes. For commodity futures, there is a wide range of contract lengths to choose from. Contracts with longer time to maturity typically have small price movements and are less liquid compared to contracts closer to maturity. The contract length that best matches the spot price movements is typically the one next to expire, i.e. the front month contract. Accordingly, we use 1 month contract lengths in our analysis, which is consistent with related studies such as Misund and Asche (2016) and Bloznelis (2018). Hedging effectiveness typically increases for cointegrated processes, such as salmon spot and futures returns, in ever longer horizons (Bloznelis, 2018). However, taking positions for hedging purposes is only sensible with appropriate forecasts of sales and feed volumes. These quantities are affected by stochastic factors such as prices and biology, which make them difficult to predict in the far future. Thus, a four-week hedging duration is chosen, which is consistent with studies such as Misund and Asche (2016).

Next, we assume that the company fully hedges other relevant exposures, most importantly currency risk, in order to focus on mitigating price risk. Previous analysis has shown that all companies in the OSLO Seafood Index hedge most of their currency exposure through currency swaps and forward contracts (Haarstad, Strypet and Strøm, 2019). Fish feed in-

[4] bu = bushel. One bushel of wheat is equivalent to 27.155 kg (CME Group, 2014).

gredients are purchased internationally in USD and EUR, so in order to remove exchange rate effects, we assume there already is a perfect currency hedge in place. We assume the same for interest rates.

Lastly, we assume a fixed transaction fee including trading and clearing when calculating transaction costs for all commodity futures contracts (Fish Pool, 2020a). The fixed fee is charged as 0.15 NOK/kg for every transaction. This will be an upper bound for the transaction costs when considering the input commodities, as agricultural commodity markets are more mature and traded on international exchanges with considerably lower costs.[5] This will to some extent compensate for extra costs associated with establishing memberships and licenses on international exchanges.

With the underlying assumptions, we define the hedged portfolio return of input and output commodities as

$$\begin{aligned}\pi(\mathbf{h}) = & \, Q^{SA}\left(s_1^{SA} - s_0^{SA}\right) - h^{SA}Q^{SA}\left(f_1^{SA} - f_0^{SA}\right) \\ & - Q^{SM}\left(s_1^{SM} - s_0^{SM}\right) + h^{SM}Q^{SM}\left(f_1^{SM} - f_0^{SM}\right) \\ & - Q^{WH}\left(s_1^{WH} - s_0^{WH}\right) + h^{WH}Q^{WH}\left(f_1^{WH} - f_0^{WH}\right) \\ & - Q^{RS}\left(s_1^{RO} - f_0^{RO}\right) + h^{RO}Q^{RO}\left(f_1^{RO} - f_0^{RO}\right),\end{aligned} \tag{3.1}$$

where superscripts SA, SM, WH, RO refer to salmon, soymeal, wheat and rapeseed oil. Q denotes the kg quantity of the commodity purchased (or sold) at the end of the hedged period. $\mathbf{h} = (h^{SA}, h^{SM}, h^{WH}, h^{RO})$ is the vector of optimal hedge ratios. s_0, f_0 denotes the initial observable spot and futures prices per kg when the hedge is set and s_1, f_1 denotes the realised spot and futures prices when the hedge is liquidated. The general subscripts $0, 1$ denote the hedge setup and liquidation times, and allow for flexible specification of different hedging horizons.

3.2 Data

In the following, we present the data used in the study and its characteristics. Spot and futures contracts price series for Atlantic salmon are denoted in NOK/kg. The spot price is a weighted average selling price based on multiple inputs, calculated on a weekly basis (Fish Pool, 2016). We convert futures prices from daily to weekly by using the final price of each week, in order to make them time consistent with the spot prices. We use the front month salmon futures prices, as discussed in Section 3.1.

Weekly price data for salmon and fish feed ingredients are obtained from Thomson Reuters Datastream for both spot and front month futures prices.[6] Price series that are denoted in bushels or tons are converted to kilograms. The feed ingredients we consider are soymeal, wheat and rapeseed oil, as discussed in Section 3.1. All observations of feed ingredient prices are converted from USD and EUR respectively, to NOK with an underlying assumption that salmon farmers already have a perfect exchange rate hedge in place, as discussed in Section 3.1. We apply fixed exchange rates of 7.0140 NOK/USD and 8.6732 NOK/EUR to obtain

[5]This is due to the fact that the nominal value of input commodities are much lower than salmon, warranting lower unit transaction costs (CME Group, 2020).

[6]Exact name and ticker for each price series obtained are listed in Appendix B.1.

prices in NOK. The exchange rates are the average NOK/USD and NOK/EUR rate over the sample period, obtained from Norges Bank (2020). The resulting price series are shown in Figure 3.1. Salmon is depicted in the top panel and soymeal, wheat and rapeseed oil in the lower panel.

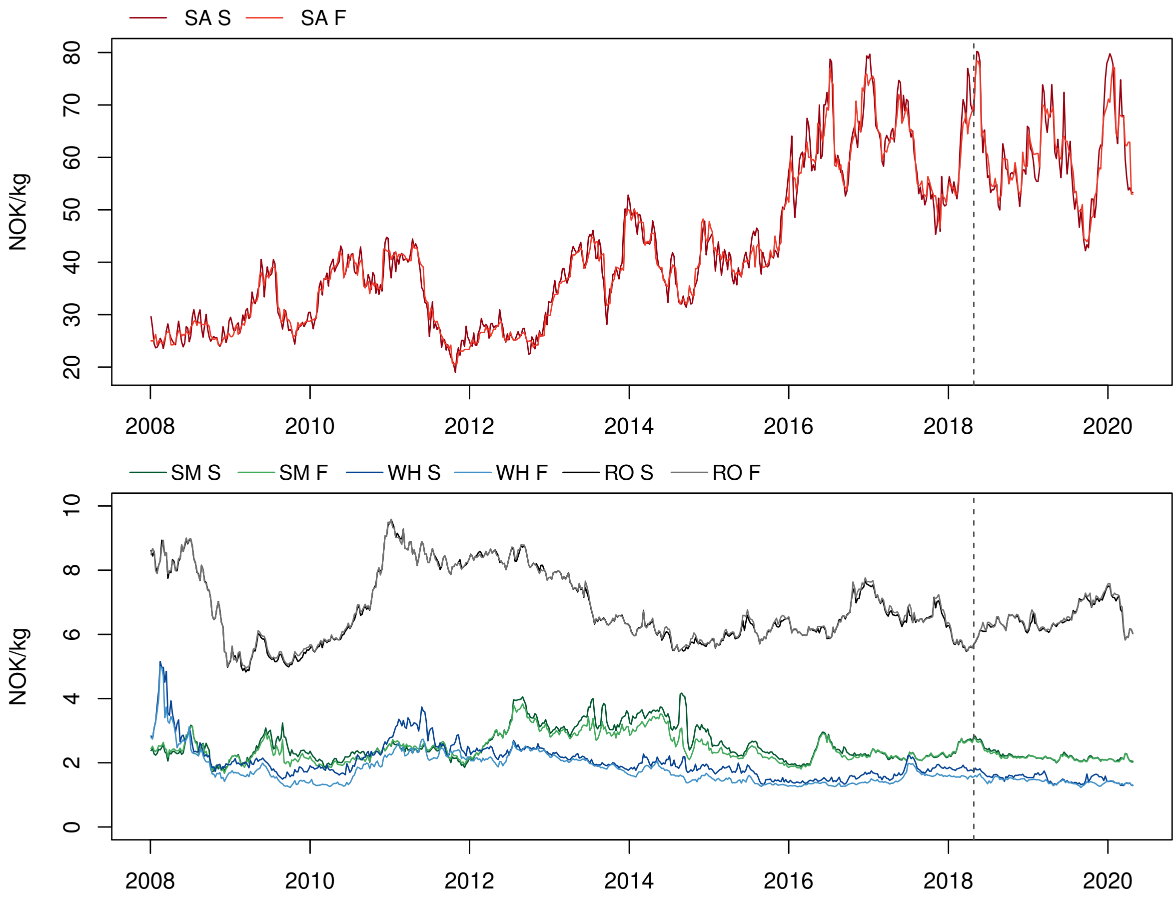

Figure 3.1: Spot and futures price series of salmon, soymeal, wheat and rapeseed oil.

Each time series consist of 643 observations collected from January 2008 to April 2020. There are no missing data points in the collected price series. For the purpose of our modelling, we use the log-transformed percentage return series which are presented in Appendix B.1, losing one observation in the process.[7] We divide the sample in two sub-samples, for estimation and hold-out samples. The 538 observations from January 2008 to April 2018 are used for model estimation, while the 104 observations between May 2018 and April 2020 form our hold-out sample. Visual inspection of the plotted log-returns suggests that most returns series feature substantial volatility clustering. The only exception is the salmon futures log-return series which seem to exhibit relatively stable volatility with a few prominent spikes evenly distributed over the sample period.

Descriptive statistics for the in-sample period are presented in Table 3.1, and show that the returns series have distributions with large variability in terms of gap between minimum and maximum observations, or standard deviation (SD). Jarque-Bera tests (JB), with null hypotheses of sample data having skewness and kurtosis matching a normal distribution,

[7] The terms *returns* and *log-returns* will be used interchangeably, referring to the same thing.

Returns series	**Mean**	**Median**	**Min**	**Max**	**SD**	**Skewn.**	**Exc.kur.**
Salmon spot	0.1563	0.0000	–18.5730	20.5376	6.2084	0.0398	0.0294
Salmon futures	0.1927	0.0000	–23.3686	22.9628	4.4423	0.0259	5.7804
Soymeal spot	0.0317	0.1939	–34.7965	24.7295	4.8265	–0.8336	8.3339
Soymeal futures	0.0225	0.1773	–29.8246	15.0282	4.5137	–0.8586	4.4427
Wheat spot	–0.0924	–0.1618	–30.9188	19.3876	5.2034	–0.3828	3.7617
Wheat futures	–0.1079	–0.1364	–32.4138	21.9919	4.3154	–0.8656	9.2657
Rapeseed oil spot	–0.0771	0.0000	–13.3531	8.6681	2.4672	–0.3719	2.7219
Rapeseed oil futures	–0.0740	0.0000	–13.3531	8.8666	2.5003	–0.2715	2.2979

Table 3.1: Descriptive statistics for in-sample weekly spot and futures returns of salmon, soymeal, wheat and rapeseed oil.

are performed and presented in Table 3.2. Results show that all bar one of the returns distributions differ significantly from the normal distribution. The distribution for salmon spot returns are close to normal, however, this series feature the largest standard deviation. These insights about the market movements that salmon farmers are exposed to highlight the need for proper risk management.

Descriptive statistics for the out-of-sample period can be found in Appendix B.1. As opposed to the in-sample period, the out-of-sample salmon and soymeal mean returns are negative, and the rapeseed mean returns are positive. Wheat mean returns are negative in both periods. This distinction between positive and negative returns will be important in discussions of results in Chapter 4.

The suitability of GARCH models depend on data assumptions as stationarity in variance. To verify this, augmented Dickey-Fuller tests (ADF) for unit roots are carried out on all returns series, with results presented in Table 3.2. All quoted test statistics are for tests with no drift and no trend.[8] Test results strongly reject the null hypotheses of a unit root for any of the returns series. These results are stable across different lag lengths.[9]

Additionally, we verify whether the returns series feature heteroscedasticity in the form of autoregressive conditional heteroscedastic (ARCH) effects. The first step is to fit the returns series to autoregressive (AR) models, with model order selection based on autocorrelation function (ACF) and partial autocorrelation function (PACF) structures, and the model selection criterion AIC. The second step before testing the residual series for ARCH effects is to ensure that the null hypothesis of no ARCH effects is not rejected due to bad fit of the AR models. Visual inspection of the ACF and PACF plots of the residuals suggest that the AR models capture the autocorrelation well. The ACF and PACF plots for salmon spot returns can be found in Appendix B.

We formally test for autocorrelation in the residuals using Ljung-Box Q test (LBQ), and find no evidence of autocorrelation.[10] Test results are presented in Table 3.2.

Finally, we can check for ARCH effects in the residuals of the fitted AR models. Formally, Engle's Lagrange multiplier (LM) test are performed on all residuals series, with results

[8]Lag length $k = 18$ is chosen for the ADF tests based on the commonly used rule of thumb by Schwert (2002), which is to choose $k = int\left(12(T/100)^{1/4}\right)$, where T denotes sample size.

[9]KPSS tests for stationarity confirm the conclusions of the ADF tests.

[10]Tests performed with 25 lags. Results are stable across a wide range of lags.

Table 3.2: In-sample test statistics.

Returns series	JB	ADF	LBQ	LM
Salmon spot	0.180	-5.61^{***}	17.78 (0.852)	44.0^{***}
Salmon futures	757.6^{***}	-5.58^{***}	21.23 (0.680)	222^{***}
Soymeal spot	1635^{***}	-5.65^{***}	16.86 (0.887)	144^{*}
Soymeal futures	514.5^{***}	-5.92^{***}	26.19 (0.398)	143^{***}
Wheat spot	334.7^{***}	-5.24^{***}	21.52 (0.664)	67.1^{***}
Wheat futures	2011^{***}	-5.12^{***}	29.81 (0.231)	140^{***}
Rapeseed oil spot	181.2^{***}	-4.77^{***}	24.24 (0.505)	85.2^{***}
Rapeseed oil futures	127.1^{***}	-4.93^{***}	28.11 (0.303)	94.1^{***}

Note: Tests applied are Jarque-Bera (JB), augmented Dickey-Fuller (ADF), Ljung-Box Q (LBQ) (p-values in parentheses) and Engle's Lagrange multiplier (LM) tests. ***, **, * denotes significance at the 1%, 5%, 10% level respectively.

presented in Table 3.2,[11] strongly rejecting the null hypotheses of no ARCH effects.[12] This is the case across a range of lag lengths, thus we can be confident that the returns series feature conditional heteroscedasticity. Tests show that the out-of-sample data feature much of the same characteristics as the in-sample data in terms of stationarity, autocorrelation and heteroscedasticity. These test statistics can be found in Appendix B.1. We conclude that GARCH models are suitable for the rest of the analysis.

3.3 Estimated models

The following section presents estimated parameters for the uni- and multivariate models described in Chapter 2, in addition to selected plots for the resulting one-ahead rolling window forecasts for conditional standard deviation and correlation.

Estimated GARCH(1,1) models

Estimated parameters and asymptotic robust standard errors (S.E.)[13] of the GARCH(1,1) models are presented in Table 3.3.[14] This is followed by the best fitted distribution, with shape parameter $\hat{\nu}$, and skew parameter $\hat{\xi}$. Furthermore, we present test statistics, critical value and p-values for the weighted ARCH LM test, the Nyblom stability test and the adjusted Pearsons goodness-of-fit test. The two-stage estimation procedure of Chapter 2 is dependant on adequately specified distributions (Patton, 2006a). It is therefore important to test the specifications of the GARCH(1,1) models, as they are the underlying structure on which the multivariate models are built.

[11] Tests performed with 12 lags.

[12] Visual inspection of Figure B.4 in Appendix B, which shows the residuals of the AR model fitted to salmon spot returns, also suggest ARCH effects are present.

[13] Method for obtaining robust standard errors are based on White (1982).

[14] ACF plot, empirical distribution plot and QQ plot for each model is attached in Appendix B.2.

Table 3.3: Estimated GARCH(1,1) models.

	Salmon S		**Salmon F**		**Soymeal S**		**Soymeal F**		**Wheat S**		**Wheat F**		**Rapeseed S**		**Rapeseed F**	
Model	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.	Est.	S.E.
$\hat{\alpha}$	0.1122	0.0422	†0.0027	0.0020	0.1124	0.0289	0.1491	0.0287	†0.1202	0.0878	0.1250	0.0061	0.0857	0.0175	0.0895	0.0214
$\hat{\beta}$	0.4968	0.1874	0.9932	0.0006	0.8278	0.0387	0.7821	0.0491	0.8472	0.1210	0.8540	0.0047	0.8680	0.0279	0.8622	0.0327
$\hat{\omega}$	15.0532	NA	0.0795	NA	1.3904	NA	1.3979	NA	0.8826	NA	0.3923	NA	0.2816	NA	0.3020	NA
	Normal		GED		t		Skewed t		Skewed t		t		t		GED	
$\hat{\nu}$	NA	NA	0.5415	0.0717	4.9027	0.9414	9.8116	3.0833	6.9182	1.4711	6.5517	1.6653	6.3924	1.3246	1.2470	0.1240
$\hat{\xi}$	NA	NA	NA	NA	NA	NA	0.8934	0.0498	1.1492	0.0747	NA	NA	NA	NA	NA	NA
$\log \mathscr{L}$	-1741		-1430		-1540		-1522		-1590		-1449		-1207		-1216	
W-LM	Stat.	p-value	Stat.	p-value	Stat.	p-value	Stat.	p-value	Stat.	p-value	Stat.	p-value	Stat.	p-value	Stat.	p-value
Lag[3]	0.0891	0.7653	0.1081	0.7423	0.292	5.89e-01	0.00418	0.9485	0.8598	0.3538	0.4509	0.5019	0.4652	0.4952	1.025	0.3112
Lag[5]	1.1043	0.7023	4.8282	0.1126	17.31	1.02e-04	0.7552	0.8067	1.1788	0.6807	1.1096	0.7007	1.187	0.6783	1.587	0.5698
Lag[7]	1.3420	0.8524	7.0233	0.0859	21.308	2.96e-05	0.8148	0.9417	1.2562	0.8687	1.6658	0.7876	1.5084	0.8197	1.858	0.7474
NS	Stat.	CV	Stat.	CV	Stat.	CV	Stat.	CV	Stat.	CV	Stat.	CV	Stat.	CV	Stat.	CV
Joint Stat.	0.3552	0.61^{***}	2.4724	1.35^{*}	1.1022	0.846^{***}	1.0179	1.07^{***}	0.4817	1.07^{***}	0.3702	0.846^{***}	0.7305	0.846^{***}	0.8421	0.846^{***}
APGoF	Stat.	p-value	Stat.	p-value	Stat.	p-value	Stat.	p-value	Stat.	p-value	Stat.	p-value	Stat.	p-value	Stat.	p-value
Group[20]	21.7	0.2993	86.24	1.53e-10	14.64	0.7452	8.022	0.9864	16.8	0.6037	4.825	0.9996	18.36	0.4987	40.59	2.74e-03
Group[30]	31	0.3656	148.65	5.04e-18	19.29	0.9137	15.829	0.9774	31.22	0.3552	15.717	0.9786	29.32	0.4483	81.96	5.96e-07
Group[40]	38.8	0.4788	206.39	1.20e-24	36.42	0.588	24.825	0.9622	48.47	0.1423	18.283	0.9981	40.88	0.3877	131.14	6.75e-12
Group[50]	49.17	0.4661	271.48	1.41e-32	38.58	0.8576	37.093	0.894	49.36	0.4587	24.268	0.9988	48.06	0.5112	177.43	1.91e-16

Note: **(1)** All coefficient estimates are significant at the 5% level, with the exception of estimates denoted with †. **(2)** W-LM: Weighted ARCH LM test for standardised GARCH residuals, the null being an adequately fitted ARCH process (Fisher and Gallagher, 2012). **(3)** NS: Nyblom stability test, CV denoting critical value, the null being constant parameter values (Nyblom, 1989). ***, **, * denotes significance at the 1%, 5%, 10% level respectively. **(4)** APGoF: Adjusted Pearson's goodness-of-fit test for p-value $(g-1)$, described in Vlaar and Palm (1993), for groups 20-50. The null being an adequately specified distribution.

Model estimates are all significant at the 5% level with exception of the salmon futures and wheat spot series, which cannot reject the null hypothesis that $\hat{\alpha} = 0$. $\hat{\alpha} \approx 0$ indicates that short term shocks have little impact on volatility, i.e. little to no volatility clustering. As mentioned in Section 3.2, this is also indicated by visual examination of the salmon futures log-return plot in Appendix B.1. However, this is not the case for the wheat spot log-returns, which show clear signs of volatility clustering by visual examination. A preliminary test using the single-hedge also show that the estimated $\hat{\alpha} = 0.1202$ outperforms $\hat{\alpha} = 0$, we therefore elect to keep the model.

The weighted ARCH LM test indicates that all models adequately capture ARCH effects at the 5% level, with exception of soymeal spot series which rejects the null of no ARCH effects for lags ≥ 5. However, performing a weighted Ljung-Box test on standardised residuals based on Fisher and Gallagher (2012), we *cannot* reject the null of no autocorrelation at the 5% level, for lags lengths ≤ 9. The weighted Ljung-Box test is also performed on the other models using squared standardised residuals, and confirms the results of the ARCH LM test. Hence, we conclude that ARCH effects are for our purposes adequately captured by the models.

The Nyblom stability test indicates that the null of constant estimated parameters cannot be rejected at the 10% level for all models, except the salmon futures and soymeal spot series. The salmon futures rejects constant parameters at the 1% level, soymeal spot at the 10% level. This is an indication of structural changes, for which a solution could be to include regime shifting models (Brooks, 2014). However, as this is not crucial for the copula approach being the focus of the thesis, we leave this for further research.

The adjusted Pearson's goodness-of-fit test is a test of whether the specified distribution adequately captures the empirical distribution. The test indicates that all conditional distributions are adequately specified, except for the salmon futures and rapeseed futures series which assume the GED distribution. For the salmon futures series the null is rejected at the 1% level, which is not surprising considering the model features no volatility clustering. However, the QQ plots attached in Appendix B.2 show that the salmon futures series are able to capture most of the quantiles except the extreme tails. QQ plot for the rapeseed futures series seems to indicate a decent specification with the exception of few extreme outliers, even though the test rejects this at the 1% level. The GED distribution was decidedly the best fitting distribution among normal, GED, Student's t and skewed t, and QQ plots seem reasonable by visual inspection. Hence, we elect to keep the distributions going forward.

To summarise, we find that the models are adequate for the goal for our methodology. One could potentially estimate more complex GARCH models to achieve higher p-values for all models. However, we consider this outside the scope of this thesis. Additionally, it will not impact the comparison and estimation of different MGARCH models since they use the same first stage models. For the sake of tractability, we focus on the copula models and keep the GARCH models simple.

Figure 3.2 shows the standardised residuals obtained from the estimated GARCH(1,1) models in Table 3.3. The matrix displays scatter plots of the resulting standardised residual pairs below the diagonal, histograms of the empirical distribution on the diagonal, and the Pearson correlation above the diagonal.

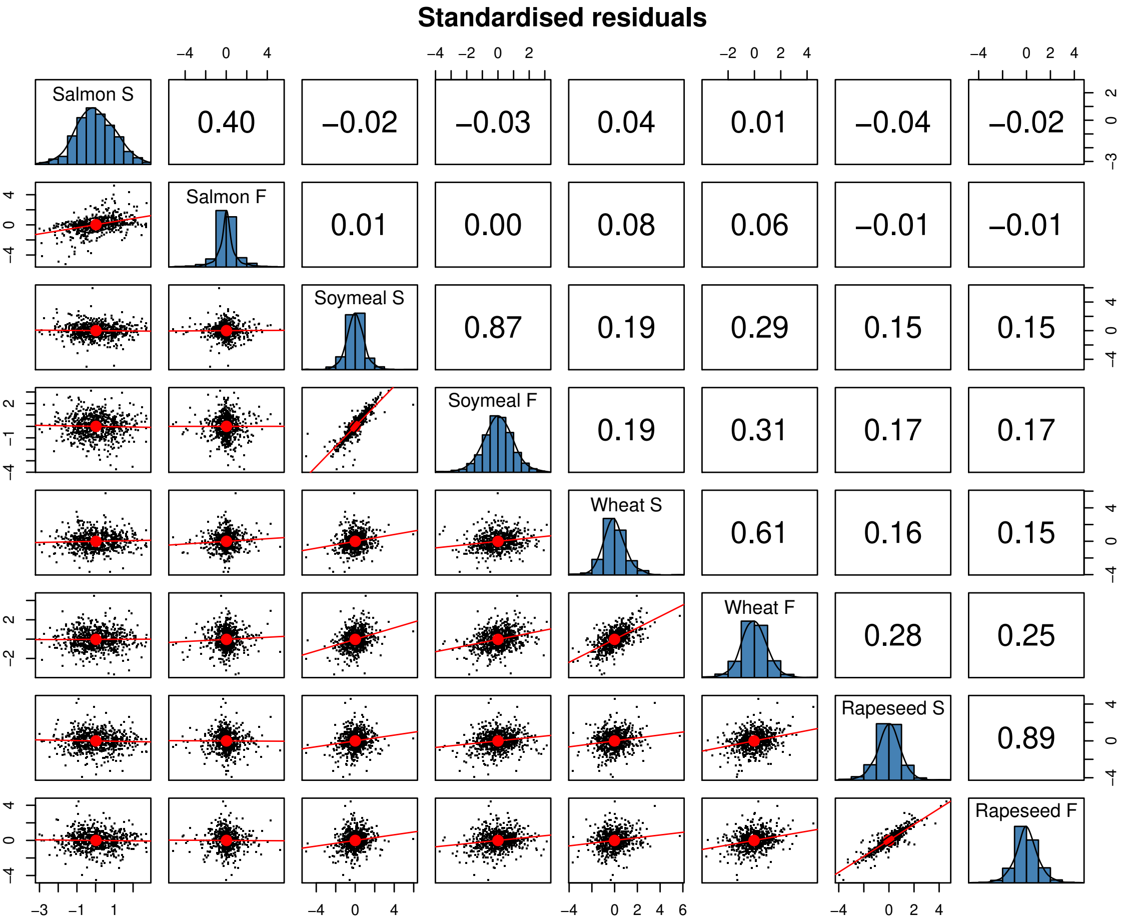

Figure 3.2: Standardised residuals obtained from estimated GARCH(1, 1) models.

Figure 3.3 presents the one-ahead rolling forecast of conditional standard deviation for the estimated GARCH(1,1) models of Table 3.3. The dashed vertical line separates the in- and out-of-sample period. The figure shows that conditional spot and futures volatility follow each other and show significant volatility clustering for all commodities except salmon. The salmon spot and futures series show substantial differences, where the salmon spot is volatile and changing frequently, with low long term volatility persistence (i.e. volatility spikes fade quickly as $(\hat{\alpha} + \hat{\beta}) << 1$).

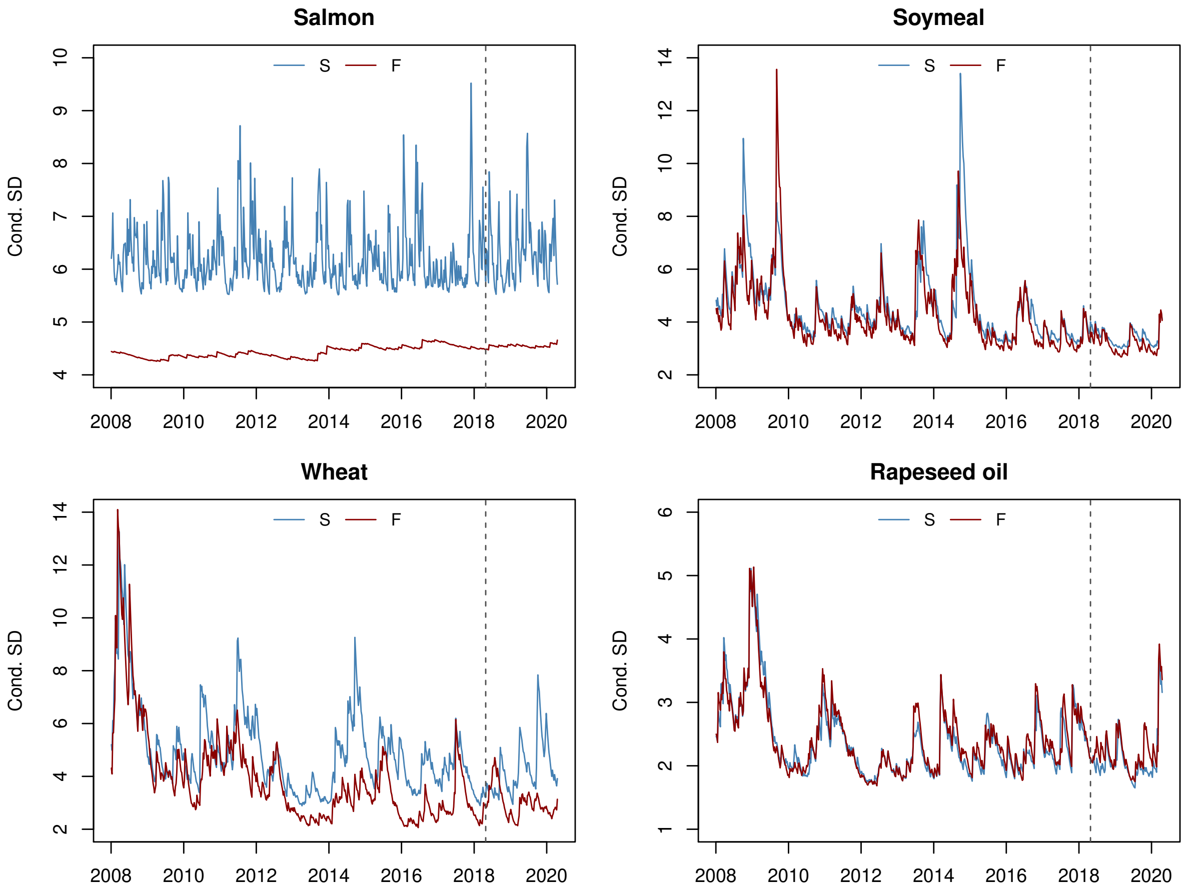

Figure 3.3: One-ahead rolling forecasts for conditional standard deviation, σ_t.

Estimated dynamic models

Table 3.4 presents the estimated parameters for the DCC type models. The DCC model assumes the multivariate normal distribution, and the C-DCC a time-varying multivariate distribution modelled by estimating a Student's t copula, as described in Section 2.3 and 2.5.

The $\tilde{\alpha}$ parameter indicate the effect of past innovations on correlation. Small $\tilde{\alpha}$ for both models imply short term shocks are less prominent in the models. Furthermore, $\tilde{\alpha}+\tilde{\beta}$ is high and implies slow decay of correlation, which is typical for financial time series (Engle, 2009). Note that $\tilde{\alpha}+\tilde{\beta}$ is slightly lower in the C-DCC model (0.8604) compared to the DCC model (0.8942). Short term effects hence have a greater impact on C-DCC correlations. $\tilde{\eta}$ is the estimated common shape parameter of the Student's t copula assumed in the C-DCC model. Notice that $\eta = 21.16$, indicating that the distribution features fat tails, although not to a large degree.[15] This is observable in Figure 3.2, where several time series pairs visually looks to be approximately normally distributed. We might therefore expect small differences when using the copula extension in the C-DCC model.

[15]A common rule of thumb is that the Student's t distribution approaches the normal distribution, the difference being negligibly small for $\eta \geq 30$ and moderately large sample sizes.

Table 3.4: Estimated parameters of the DCC(1,1) and C-DCC(1,1) model.

	DCC		C-DCC	
	Estimate	S.E.	Estimate	S.E.
$\tilde{\alpha}$	0.0229	0.0081	0.0242	0.0093
$\tilde{\beta}$	0.8713	0.0703	0.8362	0.0952
$\tilde{\eta}$	NA	NA	21.1586	4.8714
$\log \mathscr{L}$	-10 816		-10 486	

Note: All estimates are significant at the 1% level.

Table 3.5 presents the estimated parameters and asymptotic standard errors of the TVC model for each commodity. The second row indicates which of the TVC models described by Equations 2.23 is used, i.e. normal, Student's t, Gumbel or survival Gumbel. As the TVC model is more experimental, it has inherent problems with convergence when estimating the model for commodity pairs in which the long run correlation is weak and close to zero.[16] Thus we elect to solely model the spot and futures pair for each respective commodity, implying we only model the single-hedge. Additionally, we elect to use lag length $n = 1$ as this during estimation has shown to yield the most reasonable models.[17] Wu (2018) also uses lag length $n = 1$ when hedging grain sorghum and finds reasonable results.

Series	**Salmon**		**Soymeal**		**Wheat**		**Rapeseed oil**	
	Normal		Student's t		Normal		Survival Gumbel	
Model	Estimate	S.E.	Estimate	S.E.	Estimate	S.E.	Estimate	S.E.
$\overline{\alpha}$	0.1822	0.4898	-0.0448	0.0221	0.0585	0.0515	-2.3385	0.3360
$\overline{\beta}$	0.1084	1.2530	2.9605	0.1112	2.0634	1.7873	0.1588	0.0201
$\overline{\omega}$	0.7275	0.0763	0.8625	0.0856	0.1455	1.1098	1.3185	0.1120
$\overline{\nu}$	NA	NA	2.0888	0.3155	NA	NA	NA	NA
$\log \mathscr{L}$	38		566		128		551	

Table 3.5: Estimated parameters of time-varying copula models with lag length $n = 1$.

Lastly, we do not provide estimated model coefficients for the RWC model, as it uses a different bivariate copula for each variable pair at each different time step, in addition to parameters. This is impractical to show in a table, however, the results can easily be recreated by accompanying R code. Regarding rolling window size, there are limited guidance on how to select the most appropriate window. While Power et al. (2013) use a 104-week window, Misund and Asche (2016) elect to use 20- and 52-week windows in their analysis. While shorter windows are interesting, the copula approach requires moderate sample sizes to converge. By trial and error, we find the 52-week window to both be feasible and provide the best results in our analysis.

[16] Most attempts at estimation for different commodity pairs have been unsuccessful and lead to results which are difficult to interpret.

[17] In theory, larger lag lengths act as a smoothing factor, however in practice we find it difficult to estimate with lag lengths > 1 as the models tend to become unreasonable.

4 | Results

In this chapter we present the results obtained by applying the methodology of Chapter 2 to the application of a salmon producer presented in Chapter 3. First, we analyse the hedge ratios obtained by different models. Second, we examine the results for the four-week hedge. Additionally, we provide further insights by including a snapshot of a one-week hedge. Third, we analyse sensitivities to hedging horizon. Lastly, we propose a cost-effectiveness measure to elaborate on the risk-return trade-off.

We find that the multi-commodity price risk in the industry can be greatly reduced by applying a multi-hedge framework. The proposed RWC multi-hedge model reduces portfolio variance by up to 53.52% out-of-sample, and results in the best risk-return trade-off. Furthermore, extending the standard multivariate GARCH models by applying copulas result in increased performance in most cases.

Hedge ratios

We begin with an introduction of dynamic hedge ratios obtained by the models proposed in Chapter 2. Table 4.1 shows in-sample mean, minimum, maximum and standard deviation (SD) for each model and commodity.[1] For ease of comparison and consistency, all hedge ratios are presented from the same perspective, i.e. a *positive* ratio indicates a futures position *opposite* of the spot market, and a *negative* ratio indicates a futures position in the *same* direction as the spot market. In our context, a positive salmon ratio implies going *short* the futures market, as the salmon farmer is long the spot market. In the input commodities, a positive position implies going *long* the futures market, as the farmer is short the spot market.

First, we compare single-hedge ratios. All models produce relatively similar single-hedge ratio mean for the respective commodities, the largest being a 0.087 point difference in the salmon series. This is expected and consistent with previous studies (Haigh and Holt, 2000; Misund and Asche, 2016; Zhao and Goodwin, 2012). Specific differences in hedge ratios are more apparent in Figure 4.1 and 4.2, which show the plotted paths for each model. The horizontal dashed line illustrates the naïve hedge, and the vertical dashed line illustrates the separation between in- and out-of-sample. Although hedge ratio means are similar, we do find distinct differences in standard deviations. Specifically, single-hedge ratio SD for the DCC and C-DCC model are in general lower than for TVC and RWC, an indication of more stable ratios in the former models. This implies that the DCC and C-DCC single-hedge port-

[1] Statistics for out-of-sample hedge ratios show similar characteristics as in-sample and are attached in Appendix C.1.

folios require fewer adjustments between each time-step than the latter.[2] There is significantly higher SD for the TVC model in soymeal, which is a result of extreme negative spikes in Figure 4.2. Potentially, a result of extreme tail events which are more accurately captured by the copula models.[3] For the same reason, TVC is the only model with a significant negative position (-1.024), as seen by examining the minimum hedge ratios. This implies that the spot and futures price changes at these moments are estimated to be highly negatively correlated. This is rare, but not impossible, during extreme tail events (Basu and Gavin, 2017).

Table 4.1: Statistics for estimated in-sample hedge ratios.

	DCC		**C-DCC**		**TVC**	**RWC**	
	Single	Multi	Single	Multi	Single	Single	Multi
Salmon							
Mean	0.509	0.497	0.562	0.542	0.494	0.475	0.482
SD	0.077	0.077	0.115	0.109	0.128	0.170	0.227
Min	0.357	0.359	0.405	0.393	0.106	-0.016	-0.075
Max	0.949	0.937	1.231	1.204	1.428	1.056	1.341
Soymeal							
Mean	0.926	1.223	0.967	1.110	0.961	0.978	1.472
SD	0.098	0.269	0.113	0.341	0.239	0.131	1.012
Min	0.499	0.344	0.429	-0.822	-1.024	0.527	-1.656
Max	1.287	1.938	1.417	2.112	1.714	1.627	4.557
Wheat							
Mean	0.841	0.812	0.872	0.606	0.840	0.837	0.571
SD	0.183	0.432	0.207	0.461	0.190	0.192	1.914
Min	0.470	-0.510	0.409	-1.358	0.495	0.464	-4.840
Max	1.441	2.355	1.609	1.962	1.567	1.550	7.780
Rapeseed oil							
Mean	0.892	0.943	0.896	0.702	0.895	0.899	0.122
SD	0.058	0.493	0.056	0.643	0.103	0.081	2.734
Min	0.705	-0.777	0.701	-3.250	-0.002	0.655	-11.699
Max	1.134	2.414	1.135	2.410	1.146	1.134	5.298

Note: Smallest and largest standard deviation (SD) for each commodity are marked red and blue respectively.

Second, we compare multi-hedge ratios, which exhibit greater differences across the models. In general, volatile multi-hedge ratios imply that the dependency between different commodity markets is changing rapidly. Changes in dependencies lead to more favourable positions in different markets, hence more volatile ratios. As with single-hedge ratios, DCC and C-DCC models have lower SD compared to the RWC model. Interestingly, the DCC

[2] This does not directly affect transaction costs as hedges are not adjusted during the hedging horizon. If this was the case, higher SD would require more adjustments, resulting in higher transaction costs.

[3] The soymeal spot and futures pair is modelled by the Student's t TVC model, which emphasises tail dependence.

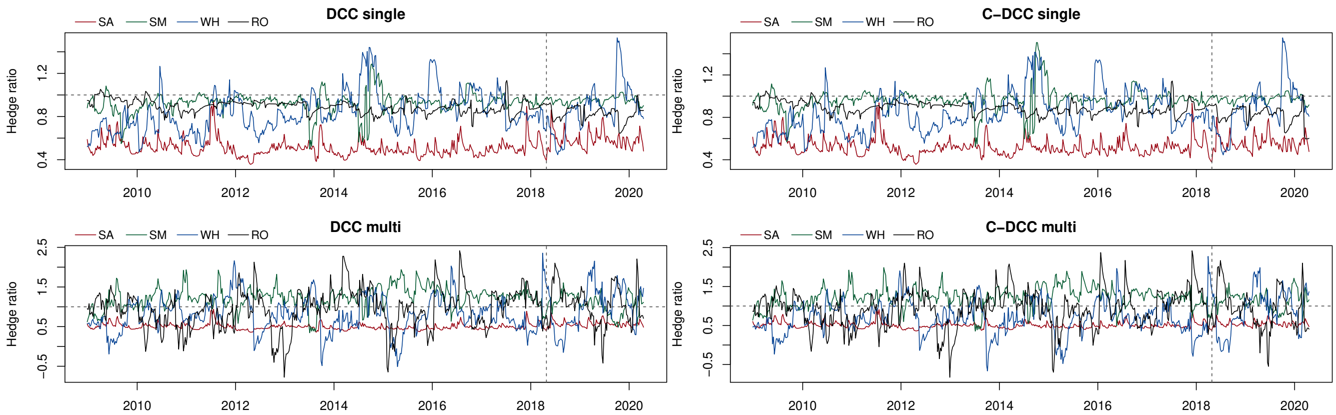

(a) DCC **(b)** C-DCC

Figure 4.1: Hedge ratios for DCC and C-DCC models.

model has significantly higher minimum hedge ratios for all input commodities, implying more conservative ratios and less extreme events being captured. The C-DCC model has noticeably more extreme values than the DCC model in general, *especially* for minimum hedge ratios. This implies that the model captures more extreme events, which result in larger negative positions. Lastly, we find that the RWC multi-hedge ratio is highly volatile for all input commodities.

Third, we compare the difference between single- and multi-hedge strategies. Most evident is the noticeable increase in SD across all models when considering the *input commodities*. This is expected, as the multi-hedge allows for both cross-hedging and speculative positions. Additionally, the volatility in ratios is significantly increased by the fact that dependency between *different* commodities is mainly driven by short term, exogenous shocks. For example, when considering soymeal and wheat, the commodities could be temporarily correlated if market participants expect a bad harvest in both markets, resulting in appreciating and correlated prices. However, if one market experience a price drop due to decreased demand, prices become negatively correlated. This rapid change in correlation has a large impact on optimal multi-hedge ratios and explains the high ratio volatility. This is not the case for the single-hedge ratio. Long run correlation between spot and futures markets on the same commodity is mainly driven by the *Law of One Price*,[4] and not temporary exogenous shocks.

Furthermore, we note that going from single- to multi-hedge barely changes the salmon hedge ratio, which is similar to findings of Anderson et al. (2017).[5] This is a crucial observation and a result of the salmon and input commodity prices being close to independent. Evidently, it is rarely possible to offset risk in the salmon price by using cross-hedges in agricultural commodities. This implies that the salmon price should be close to optimally hedged when considered alone. Consequently, any benefits of multi-hedge models that we find should come from the input commodities.

To summarise, the modelling extension of applying copulas tends to increase the stand-

[4]The Law of One Price states that market forces should align the prices of an asset over time, due to arising arbitrage opportunities.

[5]Anderson et al. (2017) finds the multi-hedge ratio of live cattle (output) approaches the single-hedge ratio, while the multi-hedge ratio of corn (input) is highly volatile in comparison.

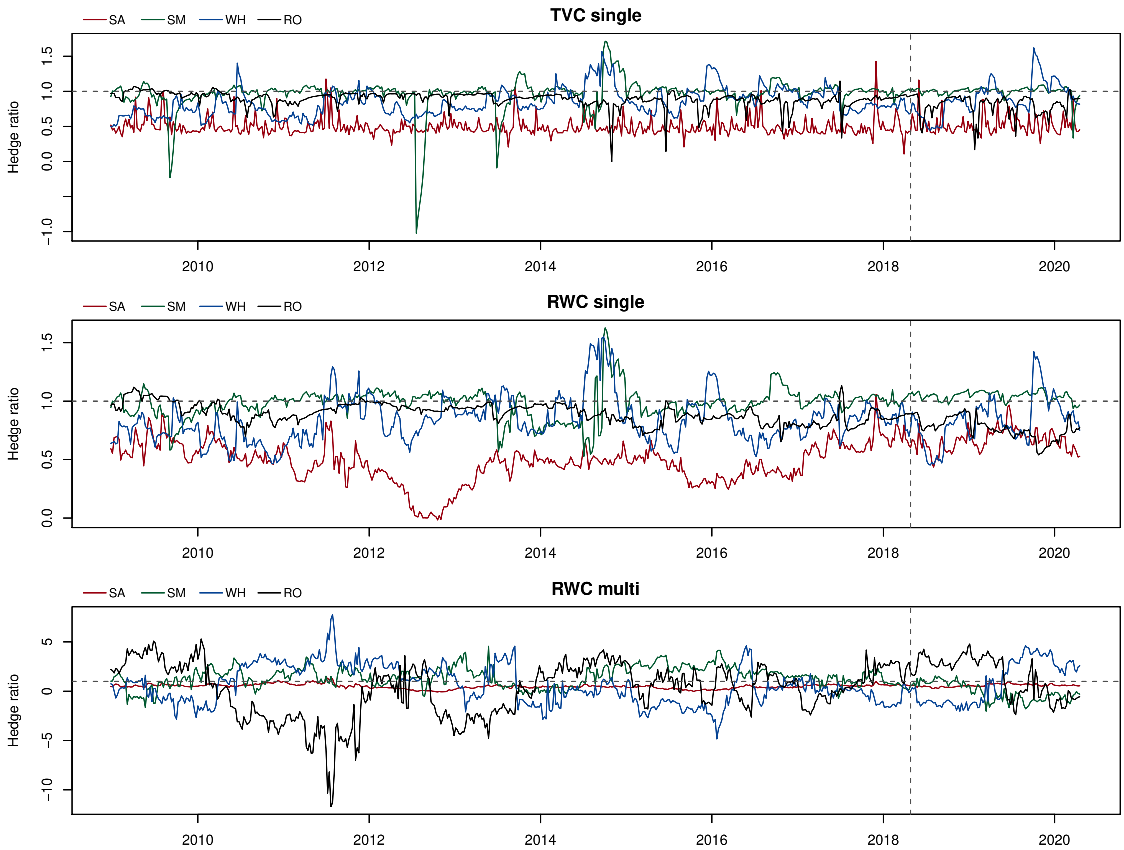

Figure 4.2: Hedge ratios for the TVC and RWC models.

ard deviation and sensitivity of single-hedge ratios, implying standard deviations in loose order of DCC < C-DCC < TVC ≈ RWC. This effect is further amplified when considering the multi-hedge ratios, which are more sensitive to small changes in correlation and variance of different commodities.

Hedging outcomes

Hedging outcomes are the result of applying Equation 3.1 to the price data of Section 3.2 and hedge ratios of Table 4.1 and C.1. The results are a distribution of hedging outcomes, which are summarised by the performance measures of Chapter 2.

Table 4.2 presents the hedging outcomes for the four-week hedging horizon. The upper half shows results for in-sample data, the lower half shows results for out-of-sample data.

Table 4.2: Hedging outcomes for a four-week hedging horizon.

	Unhedged	Naïve	DCC		C-DCC		TVC	RWC	
In-sample ($N = 538$)			Single	Multi	Single	Multi	Single	Single	Multi
Return outcomes									
Mean return	1 046	-5	571	553	571	545	606	657	962
Mean transaction cost	0	772	509	540	517	531	507	498	776
Min return	-59 589	-42 893	-34 135	-34 632	-33 570	-34 215	-33 563	-39 239	-38 243
Max return	56 700	36 674	37 455	37 450	37 439	37 434	40 048	37 495	37 459
Variance outcomes									
SD | Hedge eff. (%)	15 248	64.39%	49.04%	48.15%	49.84%	48.91%	47.70%	46.00%	45.10%
ES 05% | Reduction (%)	32 013	33.99%	26.72%	25.89%	27.33%	26.49%	26.01%	23.70%	25.31%
ES 10% | Reduction (%)	26 367	37.12%	28.90%	28.14%	29.61%	28.81%	28.61%	26.13%	27.38%
Out-of-sample ($N = 104$)									
Return outcomes									
Mean return	-2 943	-641	-1 402	-1 431	-1 480	-1 541	-1 589	-1 306	-1 177
Mean transaction cost	0	772	544	589	536	560	521	582	795
Min return	-57 975	-43 452	-45 451	-46 597	-45 658	-46 737	-47 645	-46 302	-45 589
Max return	55 327	40 964	41 844	41 907	42 234	42 321	43 742	38 590	41 361
Variance outcomes									
SD | Hedge eff. (%)	25 646	70.07%	52.24%	51.25%	50.95%	49.93%	47.66%	57.24%	53.52%
ES 05% | Reduction (%)	51 789	36.75%	28.51%	27.62%	28.16%	27.43%	26.76%	29.07%	30.16%
ES 10% | Reduction (%)	45 333	40.62%	27.14%	26.46%	26.88%	26.16%	25.31%	29.80%	28.48%

Note: Results denoted in NOK. Lowest and highest values for mean return and hedge effectiveness are marked as red and blue respectively.

First, we compare the static and dynamic strategies. For the in-sample period, all hedges yield a lower mean return (MR) than the unhedged strategy (1 046 NOK), which is to be expected when prices on average are appreciating. Interestingly, the naïve hedge outperforms all hedges in terms of hedge effectiveness (HE), reducing variance by 64.39%, significantly higher than the best dynamic model (C-DCC, 49.84%). This result is consistent with previous findings in which the naïve hedge typically outperforms other hedges in terms of minimum variance over longer horizons (Power et al., 2013; Misund and Asche, 2016).[6] This comes at a cost, as the naïve hedge has *negative* mean return, -5 NOK, in addition to the highest transaction costs, 772 NOK. The dynamic hedges yield intermediary results between the two extremes being the naïve hedge (highest HE, lowest MR), and the RWC (lowest HE, highest MR). When comparing the extremes and adjusting for transaction costs, we find that going

[6] Power et al. (2013) finds the naïve hedge to outperform other hedges in terms of minimum variance when hedging feeder cattle, while Misund and Asche (2016) finds the same for hedging in salmon futures.

from naïve to RWC, one lose 19.2 percentage points in HE (45.10% – 64.39%), but gain 963 in MR (186 – (−777)). This highlights the risk-return trade-off, which we will discuss in more detail.

Second, we compare the dynamic strategies. For the in-sample period, the RWC multi-hedge results in the highest mean return of 962 NOK, but the lowest hedge effectiveness of 45.1%. Recall from Table 4.1 that this strategy yields the most volatile hedge ratio, which also results in the highest transaction costs of 776 NOK. The most efficient models in terms of HE are the DCC (49.04%) and C-DCC (49.84%) single-hedge models, a point difference of up to 4.74% compared to the RWC multi-hedge. Nonetheless, they are among the worst performing strategies in terms of MR, both yielding 571 NOK. Recall that the DCC and C-DCC models yield relatively stable hedge ratios (low SD), and thus lower transaction costs (509 and 540 NOK). The TVC single-hedge is somewhere in the middle, with HE of 47.7% and MR of 606 NOK.

Evidently, there is an inverse relationship between variance reduction and return, as expected. When hedge effectiveness increases according to the minimum variance criteria, both downside *and* upside risk are reduced, effectively reducing the mean return. In our case, the trade-off between risk and return does not appear to be linear. We see that the RWC *multi-hedge* yields a 68.48% increase in MR and a 9.51% decrease in HE compared to the C-DCC model.[7] When considering *single-hedge* mean return compared to C-DCC, RWC has a 6.13% MR increase and a 4.29% HE decrease. TVC has a 15.06% MR increase and a 7.70% decrease. Accordingly, our results indicate a better risk-reward trade-off using the RWC-multi hedge strategy. This is also supported by the expected shortfall measure (ES). In general, we see that ES tends to be reduced as variance is reduced. It is therefore difficult to discern specific differences in ES for different models. However, we do find a significant bias toward improved ES reduction in the RWC multi-hedge model. Reduction of ES for the RWC multi-hedge (25.31%, 27.38%)[8] is greater than for the single-hedge (23.70%, 26.13%)[8]. At the same time, hedging effectiveness is greater in the single-hedge (46.00%) than the multi-hedge (45.10%). Even though HE is greater using the single-hedge, we find that the RWC multi-hedge is more efficient at reducing ES (and hence tail risk).

In Table 4.2, we observe that the dynamic multi-hedge strategies perform worse than their corresponding single-hedges in terms of HE, which might seem counter-intuitive. While the multi-hedge ratio is optimal for a given time t, it is also more time-sensitive as discussed in Section 4. The dependency estimates at time t might not hold several periods forward in time, which is why we see the multi- underperforming the single-hedge for the $t+4$ horizon, i.e. the four-week hedge. We confirm this later by examining the results of a one-week hedge.

Third, we compare the in-sample to out-of-sample results of Table 4.2. Notice that the unhedged portfolio for the out-of-sample period has a *negative* mean return of -2 943 NOK. This implies a period of overall price depreciation as described in Section 3.2, and is in contrast to the in-sample period. In this case, all hedging strategies perform better than the unhedged portfolio in terms of *both* MR and HE. The naïve hedge performs decidedly best by both measures. As the strategy with the highest hedging efficiency, it removes most of the upside during appreciating periods, yet some of the downside during depreciating periods as well. If we consider a scenario where a salmon farmer has a view on the market outlook, one could essentially optimise by using the RWC model during good periods and the naïve strategy during bad. Furthermore, we observe that the RWC model, which performed

[7] The difference becomes more extreme if we adjust for the transaction costs.

[8] Notation referring to (ES 5%, ES 10%).

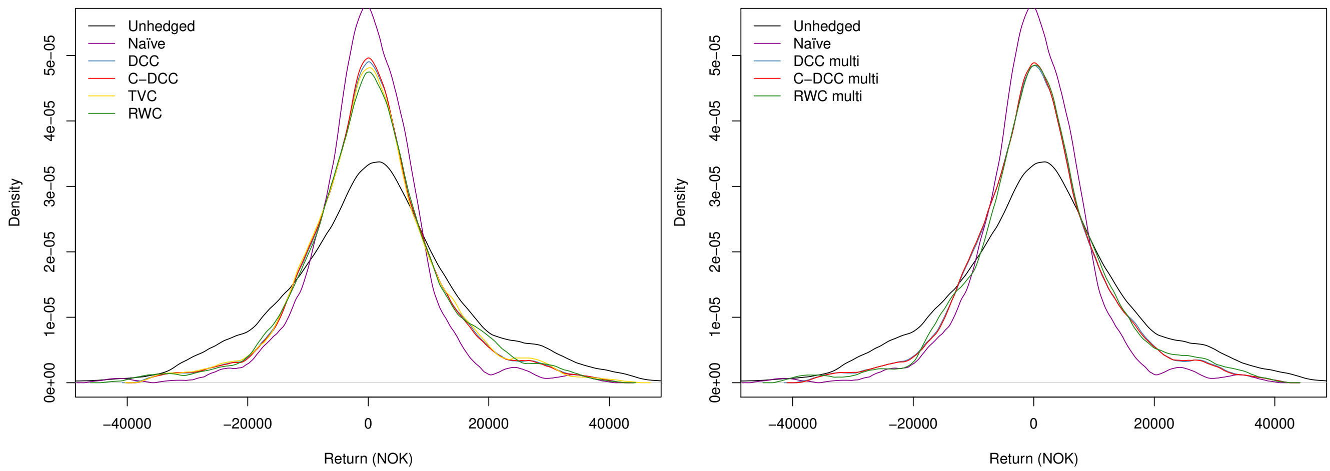

(a) Static and dynamic single-hedges.

(b) Static and dynamic multi-hedges.

Figure 4.3: In-sample return distributions for the four-week hedging horizon.

worst in terms of HE in-sample, outperforms the other dynamic models by *both* HE and MR out-of-sample. The single- and multi hedge strategies yield 57.24%, 53.52% hedge effectiveness and -1 177, -1 306 NOK mean return (-1 888, -1 972 NOK transaction cost adjusted), respectively. As a result, we find indications of the RWC model outperforming on both measures for out-of-sample. One potential reason for this might be decreasing prices during the out-of-sample period. Another reason might be related to the model specifications and differences in performance for out-of-sample data. Lastly, we find the TVC model to be the worst-performing on both HE and MR for out-of-sample. The TVC model tends to be highly sensitive to the data which the estimates are based on due to the specification of Equations 2.23, also noted by Patton (2006a). Consequently, the model does not generalise well out-of-sample.

Figure 4.3 shows the distributions of returns for different hedging strategies for the in-sample period.[9] Single-hedges are shown in the left panel and multi-hedges in the right panel. The distributions indicate where respective hedges out- and under-perform. In general, we find that the distributions for the dynamic hedges are relatively similar. However, there are piecewise differences in the tails, especially for the RWC model. This is why, for most of the models, expected shortfall (ES) tends to decrease as variance is reduced. We find the RWC multi-hedge to be the exception, for which the distribution is below the other models for most of the negative tail in Figure 4.3b. This results in lower ES as discussed previously.

Continuing, we note that the naïve hedge yields the lowest variance in returns (as returns are concentrated around the centre). Yet, it sacrifices higher returns in the approximate interval of 15 000 to 40 000 NOK. This shows the tendency of the naïve hedge to remove more of the upside than the downside (as areas below other hedge distributions are greater in the positive than the negative). Figure 4.3 reveals that while the naïve hedge yields the highest HE, most of the additional variance reduction is a result of reduction on the *upside*, and not the downside. This confirms the key take-away of our previous discussion of the risk-reward trade-off; the naïve hedge trades slight increases in variance reduction for higher reductions

[9] Out-of-sample distributions plots are attached in Appendix C.2

in return. Furthermore, we find that the dynamic models significantly reduce losses in the approximate interval of -30 000 to -15 000, but not in the -35 000 to -30 000 interval where they follow the unhedged distribution. In this segment, the naïve hedge captures some losses which are not captured by any of the dynamic models. Additionally, we see in the extreme left tail from approximately -38 000 that the naïve hedge has fatter tails than the dynamic models. Dynamic models effectively reduce more of the extreme downside losses, a desirable property when hedging. The unhedged portfolio shows high variance and heavy tails in contrast to all hedges.

Table 4.3: Hedging outcomes for a one-week hedging horizon.

	Unhedged	**Naïve**	**DCC**		**C-DCC**		**TVC**	**RWC**	
In-sample ($N = 538$)			Single	Multi	Single	Multi	Single	Single	Multi
Return outcomes									
Mean return	244	-26	132	124	135	124	146	132	192
Mean transaction	NA	772	509	540	517	531	507	498	775
Min return	-32 624	-28 666	-30 869	-30 920	-30 852	-30 894	-31 405	-31 590	-32 328
Max return	31 552	43 551	38 205	38 109	38 077	37 958	36 072	39 329	38 406
Variance outcomes									
SD | Hedge eff. (%)	8 148	0.87%	14.56%	14.64%	14,55%	14.68%	15.61%	15.97%	16.75%
ES 05% | Reduction (%)	17 930	1.66%	10.53%	10.44%	10.59%	10.47%	9.94%	11.51%	11.17%
ES 10% | Reduction (%)	14 190	-0.12%	10.46%	10.49%	10.46%	10.45%	10.18%	11.43%	11.37%
Out-of-sample ($N = 104$)									
Return outcomes									
Mean return	-628	22	-284	-298	-306	-322	-425	-267	-263
Mean transaction cost	NA	772	543	588	534	560	519	580	788
Min return	-35 401	-26 638	-27 295	-27 318	-27 366	-27 358	-27 452	-26 779	-26 678
Max return	34 431	33 031	33 560	33 608	33 603	33 648	33 559	33 384	33 301
Variance outcomes									
SD | Hedge eff. (%)	12 712	22.77%	26.27%	26.04%	25.97%	25.77%	26.00%	27.28%	26.36%
ES 05% | Reduction (%)	25 950	16.81%	20.37%	20.23%	20.49%	20.43%	18.98%	20.23%	19.54%
ES 10% | Reduction (%)	21 820	11.17%	15.72%	15.44%	15.74%	15.52%	13.98%	15.76%	15.78%

Note: Results denoted in NOK. Lowest and highest values for mean return and hedge effectiveness are marked as red and blue respectively.

Lastly, we discuss results for a one-week hedging horizon, which highlights different characteristics than the four-week hedge. As the one-week hedge has a shorter duration, it should favour dynamic hedging models. This is confirmed in Table 4.3, where dynamic models strongly outperform the static naïve hedge by all measures, both in- and out-of-sample. Additionally, we confirm that multi-hedges outperform single-hedges in terms of HE for all models in-sample, which is in line with the results of Anderson et al. (2017). Furthermore, we find they are marginally worse out-of-sample, which might stem from differences in methodology, market characteristics and sample period.[10]

To summarise, we find the RWC multi-hedge model to be the most parsimonious model. Even though it yields the lowest HE (45.10%) in-sample, it performs best on MR (962). It performs better out-of-sample, yielding the second highest HE (53.2%) and lowest MR (-1 177) among the dynamic models. Distribution plots indicate that lower HE in-sample stems from less reduction of the upside, not the downside, yielding a better risk-reward trade-off. The opposite is true for the naïve hedge. It has the highest HE, however, it performs worse on MR due to disproportionately reducing the upside. Besides, it has fatter tails on the ex-

[10] Anderson et al. (2017) uses a Monte Carlo simulation approach to approximate the multi-hedge ratio using copulas. However, they do not consider isolated in- and out-of-sample periods.

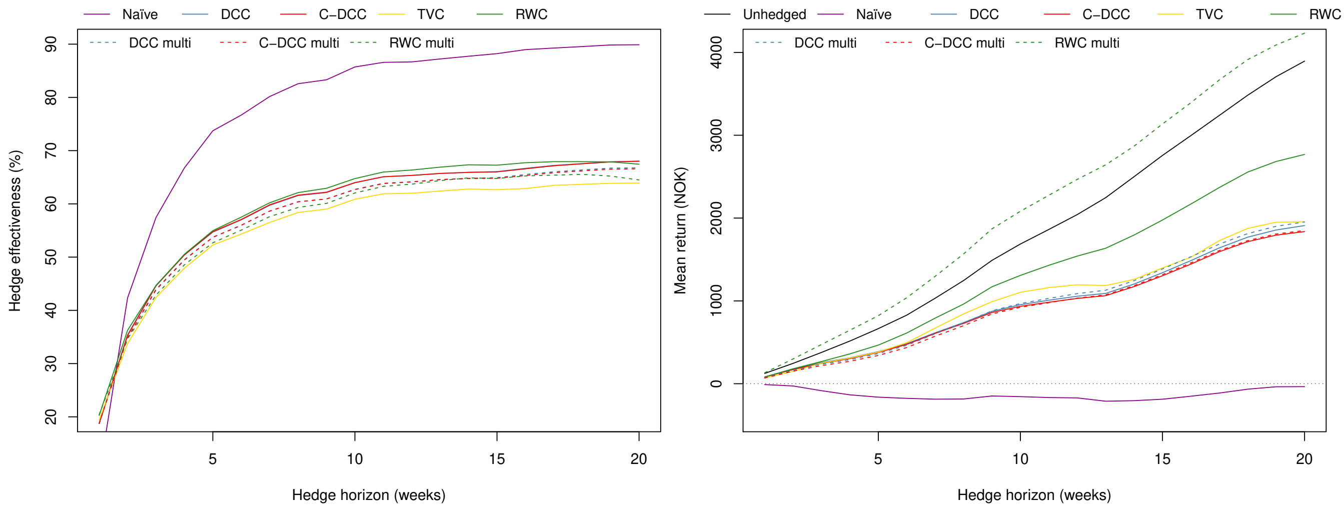

(a) Hedge effectiveness. **(b)** Mean return.

Figure 4.4: Hedging performance for different hedging horizons, full sample.

treme downside, which is undesirable when hedging. In general, we find that multi-hedging performs better in terms of HE for the one-week horizon, as expected. Furthermore, multi-hedging improves the RWC model significantly for the four-week horizon, in terms of ES and MR. Finally, we look into why shorter hedging horizons are more costly as hedging effectiveness and mean returns overall are lower, while transaction costs stay the same. There is a trade-off for salmon farmers where longer horizons are more favourable in terms of HE and costs, but requires pre-planned slaughtering volumes to a higher degree.

Sensitivity to hedging horizon

To analyse the sensitivity to different horizons, we provide results for horizons between 1 and 20 weeks for the full sample.[11] Figure 4.4a illustrates how hedging effectiveness change with the duration of the hedge. Both HE and MR increase with the hedge horizon, which is to be expected as prices are allowed to deviate more from their original values (Bloznelis, 2018). The dynamic models tend to follow each other and lead to relatively similar efficiencies, especially for hedging horizons below 4 weeks. Yet, even with a 20-week hedging horizon, there is less than 5 percentage points difference in HE between the dynamic strategies. Moreover, notice that multi-hedges perform worse than their respective single-hedges for longer horizons, as discussed previously. This is the case for the TVC model as well, which could be a reflection of the selected lag length $n = 1$.[12] The model solely captures short term dependencies and is also the worst-performing model in terms of HE for longer horizons. Further, we note that the naïve hedge outperforms on HE from the two-week horizon mark. This is partly a consequence of the over-hedging previously discussed in Section 4. Figure 4.4b clearly illustrates this, showing mean return for different horizons. When adjusting for transaction costs as shown in Figure 4.5, we find both the single- and multi-hedge RWC model outperforms other models in terms of MR, and that the gap increases with the horizon. The naïve

[11] In this section, we elect to use the full sample to focus on the sensitivity and not the specific differences between in- and out-of-sample.

[12] Modelling wise, there is limited guidance for selecting an appropriate lag length (Patton, 2006a).

hedge performs the worst, due to the low MR and high transaction costs.

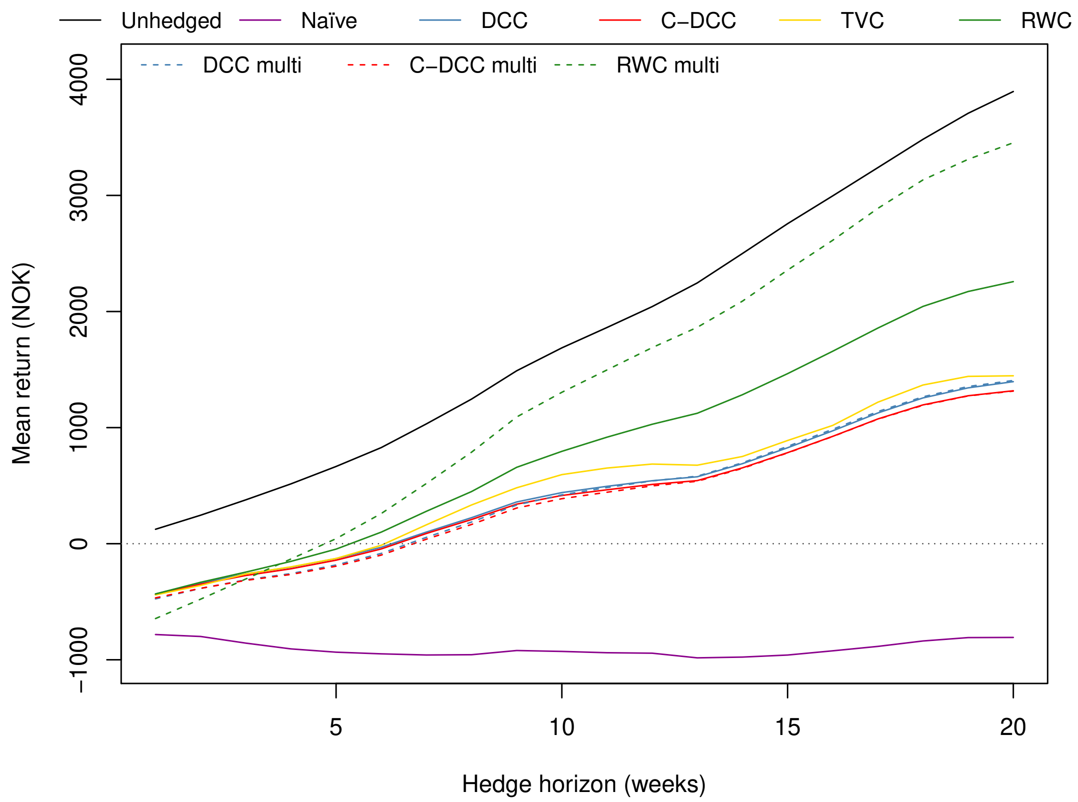

Figure 4.5: Transaction cost adjusted mean return for different hedging horizons, full sample.

Cost of hedging

In the last section we derive the implications for the trade-off between risk and return. We do this by proposing a measure of *cost-effectiveness* (CE) associated with each hedge.

First, we define *cost of hedge* (CoH) to be the difference between the transaction cost adjusted (TCA) mean return of the unhedged portfolio and the hedged. This is the return the salmon farmer historically forgoes (or gains, in case of negative cost) by using a given hedging strategy. As it also incorporates the transaction cost, we can think of it as the total cost of a given hedging strategy. A natural extension is the ratio $\frac{CoH}{HE}$, i.e. the cost per percentage of variance reduction (HE), which we denote as the cost-effectiveness[13] or CE. This measure allows us to distinguish between models in terms of the risk-return trade-off. In other words, which models yield the least costly hedge effectiveness.

[13]While the hedge effectiveness is a percentage reduction, the cost-effectiveness is a ratio of NOK to %, and thus, denoted in NOK. A lower CE is preferred, in contrast to HE.

Table 4.4: Hedging cost measures for the 4-week hedging horizon.

	Unhedged	Naïve	DCC		C-DCC		TVC	RWC	
In-sample (N = 538)			Single	Multi	Single	Multi	Single	Single	Multi
Mean return, TCA	1 046	-777	62	13	54	14	99	159	186
CoH	0	1 823	984	1 033	992	1 032	947	887	860
HE (%)	NA	64.39%	49.04%	48.15%	49.84%	48.91%	47.70%	46.00%	45.10%
CE	NA	28.31	20.07	21.45	19.90	21.10	19.86	19.29	19.07
Out-of-sample (N = 104)									
Mean return, TCA	-2 943	-1 413	-1 945	-2 020	-2 016	-2 101	-2 110	-1 888	-1 972
CoH	0	-1 530	-997	-923	-927	-842	-832	-1 055	-971
HE (%)	NA	70.07%	52.24%	51.25%	50.95%	49.93%	47.66%	57.24%	53.52%
CE	NA	-21.83	-19.09	-18.01	-18.19	-16.86	-17.47	-18.44	-18.15

Note: Results denoted in NOK. TCA: transaction cost adjusted.

Second, we compare the cost-effectiveness of models for the four-week horizon as presented in Table 4.4. Recall from the previous discussion on the risk-return trade-off that the results indicated a non-linear relationship between hedge effectiveness and mean return. If the relationship was linear, we would expect the CE to be the same for all models.[14] Meaning, an increase in HE would proportionally reduce the mean return, and thus proportionally reduce exposure to both the return downside *and* upside. This is evidently not the case as we find large differences in CE for the respective models, as shown by Table 4.4. For the in-sample period, we find the RWC multi-hedge yields the best CE (19.07), while the naïve hedge yields the worst (28.31). A point difference of 9.24. All the dynamic models are relatively similar, being in the range of 19.07 (RWC multi) to 21.45 (DCC multi). In general, we see that cost-effectiveness tends to increase with the hedge effectiveness. This implies that the reason why HE increases and variance decreases is reduction on the upside, and not the downside. The salmon farmer has to forego *more* return per unit of variance reduction, as variance reduction increases.

For the out-of-sample period, we find the reverse situation. As the period has a negative mean return, the cost is negative and is accordingly a return gain. In this case, we find the naïve hedge has the best CE (-21.83) and the C-DCC multi-hedge the worst (-16.86). This also confirms with our previous findings. However, notice that the naïve hedge outperforms the RWC single- and multi-hedge only marginally, by point differences of 3.39 and 3.68. This indicates that when the naïve hedge outperforms during periods of depreciating return, it does so only marginally. Furthermore, one should expect the RWC to perform significantly worse during depreciating returns as it performs well during appreciating returns.[15] To the contrary, we find the RWC multi-hedge model among the best performing dynamic models (-18.15) even during periods of negative return, only marginally worse than the DCC (-19.09) and C-DCC (-18.19) single-hedges.

Lastly, we examine the sensitivity of cost-effectiveness to hedging horizon. Figure 4.6 shows CE for hedging horizons between 1 and 20 weeks for in- and out-of-sample data. For in-sample we find the RWC multi-hedge model strictly outperforms other models for hedge horizons ≥ 4, and that the gap increases for longer horizons. The model is more expensive for horizons between 1 and 3 weeks, due to the increased transaction cost associated with

[14] While the previous analysis of hedging outcomes examined mean return, in this section focus on capturing the total cost of the hedge and hence, use the TCA mean return in the discussion. The conclusions of the discussion are in either case the same.

[15] To elaborate: Low reduction of mean return gains during periods of positive return should imply low reduction of mean return *losses* during periods of negative return.

the multi-hedge. This can be seen in Figure 4.5, where TCA mean return is lower for the RWC multi-hedge in horizons 1-3 compared to other models. Furthermore, the naïve hedge strictly underperforms in terms of CE for all horizons in-sample. Additionally, we find the naïve hedge to be more cost-effective for hedging horizons 1-5 for out-of-sample. This is expected since the naïve hedge has higher returns during periods of negative return. However, it should be clear from Figure 4.6b that the naïve hedge is only marginally better out-of-sample during hedging horizons 1-5, and that it *underperforms* for hedging horizons ≥ 6. Again, the RWC multi-hedge model outperforms the other models for longer horizons.

To conclude, these results confirm our previous discussions. While the naïve hedge is superior in terms of hedging effectiveness and for periods of negative returns, it tends to highly over-hedge. Meaning, when reducing variance, it predominantly does so by reducing the upside risk and potential mean return, while being most expensive in terms of transaction costs. The RWC multi-hedge model does the opposite, and tends to be the most cost-effective hedging model for longer horizons, irrespective of sample period. This is an attractive model property for salmon companies, which currently prefer being exposed to spot prices due to the fear of losing upside returns.[16]

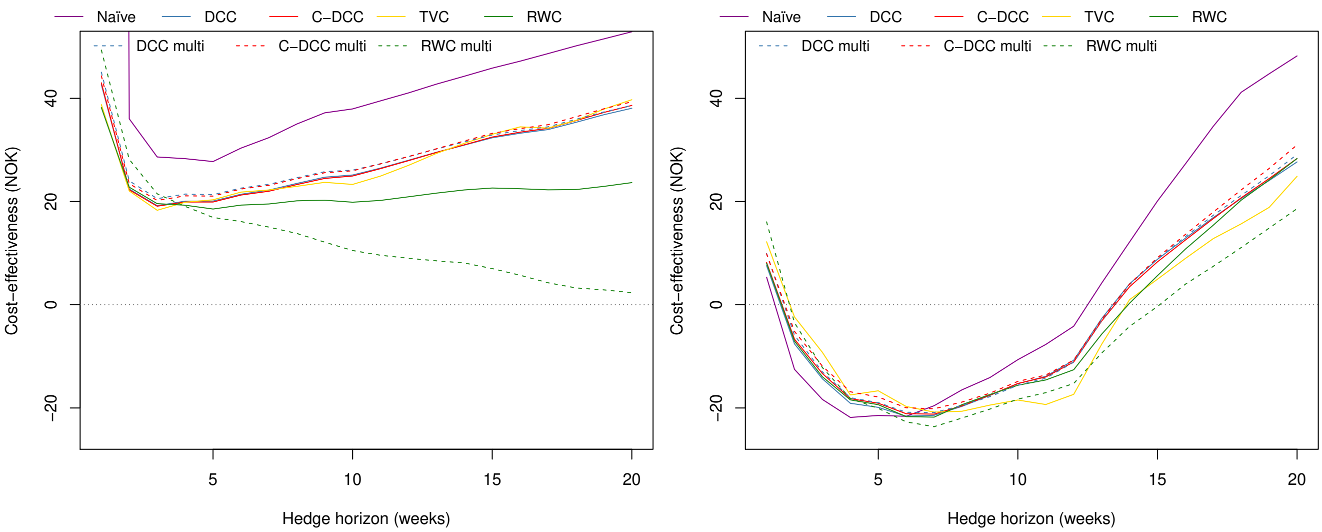

(a) In-sample. **(b)** Out-of-sample.

Figure 4.6: Cost-effectiveness for different hedging horizons.

[16]Information revealed in a phone interview with Kåre Gruven, Chief Feed Adviser at Norway Royal Salmon, 19 May 2020.

5 | Conclusion

In this study we address the price risk hedging problem for farmers of Atlantic salmon. Most industry players acknowledge the importance of price risk mitigation by engaging in trading of salmon futures or fixed-price contracts. Nonetheless, salmon farmers are exposed to risky prices not solely through their output, but also through the main production input, salmon feed. This study is the first academic contribution to hedging of joint price risk in salmon farming.

We analyse a salmon producer that partially can hedge the risk of both input and output price movements by trading in futures markets for feed ingredients and salmon. Salmon companies with integrated feed production are exposed to the same market risks, making our proposed approach applicable across multiple salmon production value chain set-ups.

Our main results can be summarised as follows. First, we find that multi-commodity price risk in the salmon farming industry can be greatly reduced by applying a state-of-the-art multi-commodity hedging framework using dynamic copula models. The proposed novel RWC multi-hedge reduces portfolio variance by 45.10% (53.52% out-of-sample) for a four-week hedging horizon. Additionally, it is the most parsimonious model sacrificing the least return per reduction of variance, and reduces expected shortfall more efficiently in comparison to other models. Although the use of the multi-hedge only improves hedging effectiveness for short hedging horizons, it tends to improve the risk-return trade-off for longer horizons.

Second, our findings indicate that the benefit of multi-hedging is a result of improved hedging of the input commodities. Using the multi-hedge, we find little changes to the optimal salmon hedging ratio. This implies it is rarely possible to offset risk in the salmon price by using cross-hedges in agricultural commodities. Furthermore, it indicates that the salmon price should be close to optimally hedged when considered alone.

Third, we find that extending the standard multivariate GARCH models by applying copulas increases hedging performance in most cases. The C-DCC model outperforms the DCC model on all measures for the in-sample four-week horizon, however, slightly underperforms out-of-sample. The largest improvement is found with the RWC model, which greatly improves the risk-return trade-off for longer hedging horizons.

Furthermore, our results show that hedging horizon greatly impacts hedging outcomes and should be considered when deciding on a hedging strategy. The hedge horizon introduces a trade-off for salmon farmers, where longer horizons are more favourable in terms of hedge effectiveness and costs, but requires pre-planning of slaughtering volumes to a higher degree.

Lastly, we propose a cost-effectiveness measure, which highlights the importance of considering the costliness of a hedge against the hedging effectiveness. The results indicate that higher hedging effectiveness comes at a disproportionate reduction of the mean return. The salmon farmer has to forego *more* return per unit of variance reduction, as variance reduction increases. The RWC model is the most efficient model in terms of cost-effectiveness for longer hedging horizons, irrespective of sample period. This is attractive for salmon companies, which currently prefer spot price exposure due to the fear of losing upside returns.

In what follows we suggest some interesting directions for future work. First is to investigate different measures for hedging effectiveness in the context of multi-commodity hedging. One alternative is lower partial moments (LPM) by Fishburn (1977), which focus more on extreme tail risk. Another measure is relative reduction in mean squared forecast error (RRMSFE) by Bloznelis (2018), which focuses on the error of forecasted expected prices. Using different hedging measures might yield different optimal hedge ratios and new industry insights.

Second, it would be interesting to explore different copula methods, like non-parametric copulas (NPC) and pair-copulas. NPC has advantages such as not assuming an elliptical dependency structure resulting in a more general copula function. Hence, it would be interesting to see if NPC potentially yields better results. NPC are usually estimated by kernel estimation, which is infeasible in higher dimensions. There are several other methods worth investigating as alternatives to the common kernel estimation. For instance, the LGDE approach using the local Gaussian correlation presented by Otneim (2016). LGDE handles higher dimensions well and is robust against dimensionality issues, modelling error, in addition to noise introduced by irrelevant parameters (Otneim and Tjøstheim, 2018). Another possibility is modelling the high-dimensional data by pair-copulas (Aas et al., 2009). Using simplified pair-copulas, one can evade the curse of dimensionality and construct higher-dimensional copulas (Nagler, Schellhase and Czado, 2017).

Another interesting direction for further research is to investigate the usage of other derivatives together with, or instead of futures. For salmon, the only financial derivative available is futures. However, Fish Pool did provide the possibility to trade Asian options earlier (Fish Pool, 2020b). The exchange has hinted that they might offer options again in the future if they can create a more liquid market. All input commodities used in our analysis have a large variety of derivatives to trade in, totalling in a vast number of possibilities for a salmon farmer seeking to reduce price risk exposure.

Salmon farmers are exposed to multiple risks, both financial and non-financial, in addition to risky prices. Additionally, several salmon companies are operating on a global scale and thus exposed to multiple exchange rates and interest rates. A natural extension is to include exchange and interest rates in the analysis, obtaining a more complete picture of the financial risk situation. However, this would also increase the number of dimensions. Another possibility is incorporation of factors such as biological shocks and optimal timing of slaughter. Finally, an interesting addition is to account for production of multiple species, which is how some of the companies in the OSLO Seafood Index operate today.

References

Aas, K. et al. (2009). ‘Pair-copula constructions of multiple dependence’. In: *Insurance: Mathematics and Economics* 44.2, pp. 182–198. DOI: 10.1016/j.insmatheco.2007.02.001.

Aas, T. S., Ytrestøyl, T. and Åsgård, T. (2019). ‘Utilization of feed resources in the production of Atlantic salmon (Salmo salar) in Norway: An update for 2016’. In: *Aquaculture Reports* 15, p. 100216. DOI: 10.1016/j.aqrep.2019.100216.

Acerbi, C. and Tasche, D. (2001). ‘On the coherence of Expected Shortfall’. In: *Journal of Banking & Finance* 26.7, pp. 1487–1503. DOI: 10.1016/S0378-4266(02)00283-2.

Anderson, D. P. et al. (2017). ‘Multi-Commodity Hedging in the Live Cattle Futures Market’. In: *Agricultural & Applied Economics Association*. DOI: 10.22004/ag.econ.258560.

Anderson, R. W. and Danthine, J.-P. (1981). ‘Cross Hedging’. In: *Journal of Political Economy* 89.6, pp. 1182–1196. DOI: 10.1086/261028.

Asche, F., Misund, B. and Oglend, A. (2016). ‘The spot-forward relationship in the Atlantic salmon market’. In: *Aquaculture Economics & Management* 20.2, pp. 222–234. DOI: 10.1080/13657305.2016.1156192.

Asche, F. et al. (2011). ‘Demand growth for Atlantic salmon: The EU and french markets’. In: *Marine Resource Economics* 26.4, pp. 255–265. DOI: 10.5950/0738-1360-26.4.255.

Basu, P. and Gavin, W. T. (2017). ‘Negative Correlation Between Stock and Futures Returns: an Unexploited Hedging Opportunity?’ In: *Bulletin of Economic Research* 69.3, pp. 209–215. DOI: 10.1111/boer.12090.

Bauwens, L., Laurent, S. and Rombouts, J. V. (2006). ‘Multivariate GARCH models: A survey’. In: *Journal of Applied Econometrics* 21.1, pp. 79–109. DOI: 10.1002/jae.842.

Berge, A. (2019). *Her er fallgruvene for oppdrett i RAS-anlegg*. URL: https://ilaks.no/her-er-fallgruvene-for-oppdrett-i-ras-anlegg/ (visited on 14/10/2019).

BioMar Group (2018). *Sustainability Report 2018*. URL: https://www.biomar.fr/globalassets/.global/pdf-files-_en/biomar-group-sustainability-report-2018.pdf (visited on 20/09/2019).

— (2019). *Annual Report 2018*. Aarhus, Denmark. URL: https://www.biomar.fr/globalassets/.global/pdf-files-_en/biomar-group-annual-report-2018.pdf (visited on 20/09/2019).

Bloznelis, D. (2018). ‘Hedging salmon price risk’. In: *Aquaculture Economics & Management* 22.2, pp. 168–191. DOI: 10.1080/13657305.2017.1409292.

Brækkan, E. H. (2014). ‘Disentangling supply and demand shifts: the impacts on world salmon price’. In: *Applied Economics* 46.32, pp. 3942–3953. DOI: 10.1080/00036846.2014.948674.

Brechmann, E. C. and Schepsmeier, U. (2013). ‘Modeling dependence with C- and D-vine copulas: The R package CDVine’. In: *Journal of Statistical Software* 52.3, pp. 1–27. DOI: 10.18637/jss.v052.i03.

Brooks, C. (2014). *Introductory Econometrics for Finance*. 3rd. Cambridge University Press (CUP).

Butterworth, D. and Holmes, P. (2001). 'The hedging effectiveness of stock index futures: evidence for the FTSE-100 and FTSE-mid250 indexes traded in the UK'. In: *Applied Financial Economics* 11.1, pp. 57–68. DOI: 10.1080/09603100150210264.

Caporin, M. and McAleer, M. (2014). 'Robust ranking of multivariate GARCH models by problem dimension'. In: *Computational Statistics and Data Analysis* 76.June, pp. 172–185. DOI: 10.1016/j.csda.2012.05.012.

Cherubini, U. (2004). *Copula methods in finance*. Wiley finance series. Chichester: Wiley.

CME Group (2014). *Agricultural Commodity Metric Conversion Guide (AC225/400/0214)*. Chicago, Illinois. URL: https://www.cmegroup.com/trading/agricultural/files/AC-225_MetricGuide.pdf (visited on 04/05/2020).

— (2020). *Exchange Fees for Clearing & Trading*. URL: https://www.cmegroup.com/company/clearing-fees.html (visited on 29/05/2020).

Ederington, L. H. (1979). 'The Hedging Performance of the New Futures Markets'. In: *The Journal of Finance* 34.1, p. 157. DOI: 10.2307/2327150.

Engle, R. (2002). 'Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models'. In: *Journal of Business and Economic Statistics* 20.3, pp. 339–350. DOI: 10.1198/073500102288618487.

— (2009). *Anticipating Correlations: A New Paradigm for Risk Management*. The Econometric and Tinbergen Institutes Lectures. US: Princeton University Press.

EY (2019). *The Norwegian Aquaculture Analysis 2019*. URL: https://www.ey.com/Publication/vwLUAssets/Norwegian_Aquaculture_Analysis_2019/%5C$FILE/The%5C%20Norwegian%5C%20Aquaculture%5C%20Analysis_2019.pdf (visited on 19/05/2020).

Fackler, P. L. and McNew, K. P. (1993). 'Multiproduct Hedging: Theory, Estimation, and an Application'. In: *Review of Agricultural Economics* 15.3, p. 521. DOI: 10.2307/1349486.

Fama, E. F. (1990). 'Stock Returns, Expected Returns, and Real Activity'. In: *The Journal of Finance* 45.4, p. 1089. DOI: 10.2307/2328716.

Fama, E. F. and French, K. R. (1992). 'The Cross-Section of Expected Stock Returns'. In: *The Journal of Finance* 47.2, pp. 427–465. DOI: 10.1111/j.1540-6261.1992.tb04398.x.

Fernández, C. and Steel, M. F. (1998). 'On bayesian modeling of fat tails and skewness'. In: *Journal of the American Statistical Association* 93.441, pp. 359–371. DOI: 10.1080/01621459.1998.10474117.

Fish Pool (2016). *Fish Pool Index™*. URL: http://fishpool.eu/price-information/spot-prices/fish-pool-index/ (visited on 03/11/2019).

— (2020a). *Fee list*. URL: http://fishpool.eu/trading/fee-list/ (visited on 25/05/2020).

— (2020b). *Options*. URL: http://fishpool.eu/products/options/ (visited on 21/06/2020).

Fishburn, P. C. (1977). 'Mean-Risk Analysis with Risk Associated with Below-Target Returns'. In: *American Economic Review* 67.2, pp. 116–126. URL: https://www.jstor.org/stable/1807225.

Fisher, T. J. and Gallagher, C. M. (2012). 'New weighted portmanteau statistics for time series goodness of fit testing'. In: *Journal of the American Statistical Association* 107.498, pp. 777–787. DOI: 10.1080/01621459.2012.688465.

Fiskeridirektoratet (2020). *Kapasitetsjustering / trafikklyssystem 2017-2018*. URL: https://fiskeridir.no/Akvakultur/Tildeling-og-tillatelser/Kapasitetsjustering-trafikklyssystemet/Kapasitetsjustering-trafikklyssystem-2017-2018 (visited on 23/01/2020).

Franken, J. R. and Parcell, J. L. (2011). 'Cross-hedging fishmeal: Exploring corn and soybean meal futures contracts'. In: *Aquaculture Economics and Management* 15.1, pp. 71–81. DOI: 10.1080/13657305.2011.549166.

French, K. R. and Fama, E. F. (1989). 'Business Conditions and Expected Returns on Stocks and Bonds'. In: *Journal of Financial Economics* 25, pp. 23–49. DOI: 10.1016/0304-405X(89)90095-0.

Genest, C., Ghoudi, K. and Rivest, L. p. (1995). 'A semiparametric estimation procedure of dependence parameters in multivariate families of distributions'. In: *Biometrika* 82.3, pp. 543–552. DOI: 10.1093/biomet/82.3.543.

Ghalanos, A. (2019a). *rmgarch: Multivariate GARCH models*. R package version 1.3-6. URL: https://cran.r-project.org/web/packages/rmgarch/rmgarch.pdf.

— (2019b). 'The rmgarch models : Background and properties (Version 1.3-0)'. URL: https://cran.r-project.org/web/packages/rmgarch/vignettes/The_rmgarch_models.pdf.

— (2020). *rugarch: Univariate GARCH models*. R package version 1.4-2. URL: https://cran.r-project.org/package=rugarch.

Graham, J. R. and Smith, C. W. (1999). 'Tax incentives to hedge'. In: *Journal of Finance* 54.6, pp. 2241–2262. DOI: 10.1111/0022-1082.00187.

Grieg Seafood (2019). *Annual Report 2018*. URL: https://www.griegseafood.no/wp-content/uploads/2019/04/Grieg_A%5C%CC%5C%8Arsrapport__Komplett_Enkeltsider_110419.pdf (visited on 16/02/2020).

Haarstad, A. H., Strypet, K. and Strøm, E. (2019). 'Cross-hedging salmon feed through market-based instruments: A structural modelling approach'. Project thesis. Trondheim: NTNU.

Haigh, M. S. and Holt, M. T. (2000). 'Hedging multiple price uncertainty in international grain trade'. In: *American Journal of Agricultural Economics* 82.4, pp. 881–896. DOI: 10.1111/0002-9092.00088.

Hansen, P. R. and Lunde, A. (2005). 'A forecast comparison of volatility models: does anything beat a GARCH(1,1)?' In: *Journal of Applied Econometrics* 20.7, pp. 873–889. URL: https://www.jstor.org/stable/25146403.

Hofert, M. et al. (2020). *copula: Multivariate Dependence with Copulas*. R package version 1.0-0. URL: https://cran.r-project.org/package=copula.

Hovland, K. M., Hopland, S. and Solheimsnes, P. (2019). *Tusenvis av tonn med laks er død: Disse rammes av algeinvasjonen*. URL: https://e24.no/naeringsliv/i/9v84l9/tusenvis-av-tonn-med-laks-er-doed-disse-rammes-av-algeinvasjonen (visited on 22/01/2020).

Hyndman, R. et al. (2020). *forecast: Forecasting functions for time series and linear models*. R package version 8.12. URL: https://cloud.r-project.org/package=forecast.

Jensen, P. M. (2019). *Hva er begrensningene for videre vekst i oppdrett? Dette mener Aqua Nor-deltagere*. URL: https://www.kyst.no/article/hva-er-begrensningene-for-videre-vekst-i-oppdrett-dette-mener-aqua-nor-deltagere/ (visited on 17/02/2020).

Joe, H. (1997). *Multivariate models and dependence concepts*. Vol. 73. Monographs on statistics and applied probability. London: Chapman & Hall.

Jondeau, E. and Rockinger, M. (2006). 'The Copula-GARCH model of conditional dependencies: An international stock market application'. In: *Journal of International Money and Finance* 25.5, pp. 827–853. DOI: 10.1016/j.jimonfin.2006.04.007.

King, R. (2012). *skewt: The Skewed Student-t Distribution*. R package version 0.1. URL: https://cran.r-project.org/package=skewt.

Knudsen, C. (2019). *Laveste laksepris på 4 år: – Bunnen er nådd*. URL: https://e24.no/boers-og-finans/i/EWEj73/laveste-laksepris-paa-4-aar-bunnen-er-naadd (visited on 20/11/2019).

Lerøy Seafood Group (2019). *Annual Report 2018*. URL: https://www.leroyseafood.com/globalassets/02-documents/rapporter/arsrapporter/arsrapport-2018 (visited on 29/09/2019).

Misund, B. and Asche, F. (2016). 'Hedging efficiency of Atlantic salmon futures'. In: *Aquaculture Economics and Management* 20.4, pp. 368–381. DOI: 10.1080/13657305.2016.1212123.

Mowi (2019). *Salmon Farming Industry Handbook 2019*. URL: https://corpsite.azureedge.net/corpsite/wp-content/uploads/2019/06/Salmon-Industry-Handbook-2019.pdf (visited on 02/10/2019).

— (2020). *Integrated Annual Report 2019*. URL: https://corpsite.azureedge.net/corpsite/wp-content/uploads/2020/03/Mowi_Annual_Report_2019.pdf (visited on 01/05/2020).

Nagler, T. and Czado, C. (2016). 'Evading the curse of dimensionality in nonparametric density estimation with simplified vine copulas'. In: *Journal of Multivariate Analysis* 151, pp. 69–89. DOI: 10.1016/j.jmva.2016.07.003.

Nagler, T., Schellhase, C. and Czado, C. (2017). 'Nonparametric estimation of simplified vine copula models: Comparison of methods'. In: *Dependence Modeling* 5.1, pp. 99–120. DOI: 10.1515/demo-2017-0007. arXiv: 1701.00845.

Nagler, T. et al. (2019). *VineCopula: Statistical Inference of Vine Copulas*. R package version 2.3.0. URL: https://cran.r-project.org/package=VineCopula.

Narasimhan, B. et al. (2019). *cubature: Adaptive Multivariate Integration over Hypercubes*. R package version 2.0.4. URL: https://cran.r-project.org/package=cubature.

Neuwirth, E. (2014). *RColorBrewer: ColorBrewer Palettes*. R package version 1.1-2. URL: https://cran.r-project.org/package=RColorBrewer.

Norges Bank (2020). *Valutakurser*. URL: https://www.norges-bank.no/tema/Statistikk/Valutakurser/ (visited on 24/05/2020).

Norway Royal Salmon (2020). *Årsrapport 2019*. URL: https://norwayroyalsalmon.com/files/9/%C3%85rsrapport%202019.pdf (visited on 01/05/2020).

Nyblom, J. (1989). 'Testing for the constancy of parameters over time'. In: *Journal of the American Statistical Association* 84.405, pp. 223–230. DOI: 10.1080/01621459.1989.10478759.

Oglend, A. (2013). 'Recent Trends in Salmon Price Volatility'. In: *Aquaculture Economics and Management* 17.3, pp. 281–299. DOI: 10.1080/13657305.2013.812155.

Ooi, H., Microsoft and Weston, S. (2019). *foreach: Provides Foreach Looping Construct*. R package version 1.5.0. URL: https://cran.r-project.org/package=foreach.

Otneim, H. (2016). 'Multivariate and conditional density estimation using local Gaussian approximations'. Dissertation for the degree of Philosophiae Doctor (PhD). University of Bergen. URL: https://hdl.handle.net/1956/15333.

Otneim, H. and Tjøstheim, D. (2018). 'Conditional density estimation using the local Gaussian correlation'. In: *Statistics and Computing* 28.2, pp. 303–321. DOI: 10.1007/s11222-017-9732-z.

Parcell, J. and Franken, J. (2011). *Agricultural Commodity Futures Contract Specifications*. Columbia: Department of Agricultural and Applied Economics. URL: https://extension2.missouri.edu/G601.

Patton, A. J. (2006a). 'Estimation of multivariate models for time series of possibly different lengths'. In: *Journal of Applied Econometrics* 21.2, pp. 147–173. DOI: 10.1002/jae.865.

— (2006b). 'Modelling asymmetric exchange rate dependence*'. In: *International Economic Review* 47.2, pp. 527–556. DOI: 10.1111/j.1468-2354.2006.00387.x.

— (2012). 'A review of copula models for economic time series'. In: *Journal of Multivariate Analysis* 110, pp. 4–18. DOI: 10.1016/j.jmva.2012.02.021.

— (2020a). *Andrew Patton's Matlab code page*. URL: https://public.econ.duke.edu/%7B~%7Dap172/code.html (visited on 26/03/2020).

Patton, A. J. (2020b). *Andrew Patton's research page.* URL: https://public.econ.duke.edu/%7B~%7Dap172/research.html (visited on 08/03/2020).

Plate, T. and Heiberger, R. (2016). *abind: Combine Multidimensional Arrays.* R package version 1.4-5. URL: https://cran.r-project.org/package=abind.

Power, G. J. and Vedenov, D. (2008). 'The Shape of the Optimal Hedge Ratio: Modeling Joint Spot-Futures Prices using an Empirical Copula-GARCH Model'. In: *Proceedings of the NCCC-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management.* URL: https://legacy.farmdoc.illinois.edu/nccc134/conf_2008/pdf/confp11-08.pdf.

— (2009). 'Dealing with downside risk in a multi-commodity setting: A case for a "Texas hedge"?' In: *Journal of Futures Markets* 30.3, pp. 290–304. DOI: 10.1002/fut.20411.

Power, G. J. et al. (2013). 'Market volatility and the dynamic hedging of multi-commodity price risk'. In: *Applied Economics* 45.27, pp. 3891–3903. DOI: 10.1080/00036846.2012.736942.

Qiu, D. (2015). *aTSA: Alternative Time Series Analysis.* R package version 3.1.2. URL: https://cran.r-project.org/package=aTSA.

Revelle, W. (2019). *psych: Procedures for Psychological, Psychometric, and Personality Research.* R package version 1.9.12. Northwestern University. Evanston, Illinois. URL: https://cran.r-project.org/package=psych.

Rosenberg, J. V. and Schuermann, T. (2006). 'A general approach to integrated risk management with skewed, fat-tailed risks'. In: *Journal of Financial Economics* 79.3, pp. 569–614. DOI: 10.1016/j.jfineco.2005.03.001.

Ryan, J. A. et al. (2020). *quantmod: Quantitative Financial Modelling Framework.* R package version 0.4.17. URL: https://cran.r-project.org/package=quantmod.

SalMar (2020). *Annual report 2019.* URL: https://ml-eu.globenewswire.com/Resource/Download/74e5d277-5302-41ee-af23-4a7785858807 (visited on 04/05/2020).

Salmon Facts (2016). *How do farmed salmon affect CO2 emissions?* URL: https://salmonfacts.com/salmon-and-environment/how-does-farmed-salmon-affect-co2-emissions/ (visited on 22/01/2020).

Schmidbauer, H., Roesch, A. and Tunalioglu, V. S. (2016). *mgarchBEKK: Simulating, Estimating and Diagnosing MGARCH (BEKK and mGJR) Processes.* R package version 0.0.2. URL: https://cran.r-project.org/package=mgarchBEKK.

Schwert, G. W. (2002). 'Tests for Unit Roots: A Monte Carlo Investigation'. In: *Journal of Business & Economic Statistics* 20.1, pp. 5–17. DOI: 10.1198/073500102753410354.

Seafish (2012). *Responsible Sourcing Guide: Farmed Atlantic salmon.* URL: https://www.seafish.org/media/publications/SeafishResponsibleSourcingGuide_Atlanticsalmon_201208.pdf (visited on 14/10/2019).

Sklar, A. (1959). 'Fonctions de répartition à n dimensions et leurs marges'. In: *Publications de l'Institut Statistique de l'Université de Paris* 8, pp. 229–231.

Smith, C. W. and Stulz, R. M. (1985). 'The Determinants of Firms' Hedging Policies'. In: *The Journal of Financial and Quantitative Analysis* 20.4, pp. 391–405. DOI: 10.2307/2330757.

Thyholdt, S. B. (2014). 'The importance of temperature in farmed salmon growth: Regional growth functions for Norwegian farmed salmon'. In: *Aquaculture Economics & Management* 18.2, pp. 189–204. DOI: 10.1080/13657305.2014.903310.

Toft, K. B. (1996). 'On the Mean-Variance Tradeoff in Option Replication with Transactions Costs'. In: *The Journal of Financial and Quantitative Analysis* 31.2, pp. 233–263. DOI: 10.2307/2331181.

Trapletti, A. and Hornik, K. (2019). *tseries: Time Series Analysis and Computational Finance*. R package version 0.10-47. URL: https://cran.r-project.org/package=tseries.

Tse, Y. K. and Tsui, A. K. (2002). 'A multivariate generalized autoregressive conditional heteroscedasticity model with time-varying correlations'. In: *Journal of Business and Economic Statistics* 20.3, pp. 351–362. DOI: 10.1198/073500102288618496.

Vlaar, P. J. and Palm, F. C. (1993). 'The message in weekly exchange rates in the european monetary system: Mean reversion, conditional heteroscedasticity, and jumps'. In: *Journal of Business and Economic Statistics* 11.3, pp. 351–360. DOI: 10.1080/07350015.1993.10509963.

Vukina, T. and Anderson, J. L. (1993). 'A State-Space Forecasting Approach to Optimal Intertemporal Cross-Hedging'. In: *American Journal of Agricultural Economics* 75.2, pp. 416–424. DOI: 10.2307/1242926.

White, H. (1982). 'Maximum Likelihood Estimation of Misspecified Models'. In: *Econometrica* 50.1, p. 1. DOI: 10.2307/1912526.

Wickham, H. et al. (2019). *readxl: Read Excel Files*. R package version 1.3.1. URL: https://cran.r-project.org/package=readxl.

Wu, S. (2018). 'Cross-hedging with Agricultural Commodities: A Copula-based GARCH Model'. Raleigh. URL: http://shenanwu.wordpress.ncsu.edu/files/2018/05/Cross-hedging-with-Agricultural-Commodities2.pdf.

Wuertz, D. et al. (2020). *fGarch: Rmetrics - Autoregressive Conditional Heteroskedastic Modelling*. R package version 3042.83.2. URL: https://cran.r-project.org/package=fGarch.

Zhao, J. and Goodwin, B. K. (2012). 'Dynamic Cross-Hedge Ratios: An Application of Copula Models'. In: *Agricultural & Applied Economics Association's 2012 AAEA Annual Meeting*. Seattle, Washington. URL: https://ageconsearch.umn.edu/record/124610/files/AAEA%202012%20Dynamic%20Cross-Hedge%20Ratios_%20Zhao_%201.pdf.

A Methodology

A.1 Copula Models

Table A.1: Copula models overview of some of the most common copula models.

	Parameter(s)	Parameter space	Independence	Pos & Neg dependency	Lower tail dependency	Upper tail dependency
Normal	ρ	$(-1,1)$	0	Yes	0	0
Student's t	ρ,ν	$(-1,1)\times(2,\infty)$	$(0,\infty)$	Yes	$g_T(\rho,\nu)$	$g_T(\rho,\nu)$
Clayton	γ	$(0,\infty)$	0	No †	$2^{-\frac{1}{\gamma}}$	0
Rotated Clayton	γ	$(0,\infty)$	0	No †	0	$2^{-\frac{1}{\gamma}}$
Gumbel	γ	$(1,\infty)$	1	No	0	$2-2^{\frac{1}{\gamma}}$
Rotated Gumbel	γ	$(0,\infty)$	1	No	$2-2^{\frac{1}{\gamma}}$	0
Frank	γ	$(-\infty,\infty)$	0	Yes	0	0

Note: The independence column show the values that lead to the independence copula.
† Clayton and the rotated versions of it allow for negative dependence for $\gamma\epsilon(-1,0)$, which is different from the positive dependence case, $\gamma>0$. See Patton (2012, p.63) for more details regarding the table.

A.2 Copula densities

Frank copula density:

$$c_F(u_1,u_2;\rho)=\frac{\rho\eta e^{-\rho(u_1+u_2)}}{\left(\eta-(1-e^{\rho u_1})(1-e^{-\rho u_2})\right)^2}, \tag{A.1}$$

for $0\le\rho<\infty$) where $\eta=1-e^{-\rho}$ (Joe, 1997, p.141).

Joe copula density:

$$c_J(u_1,u_2;\rho)=(u_1^\rho+u_2^\rho-u_1^\rho u_2^\rho)^{-2+\frac{1}{\rho}}u_1^{\rho-1}u_2^{\rho-1}(\rho-1+u_1^\rho u_2^\rho-u_1^\rho u_2^\rho), \tag{A.2}$$

for $1\le\rho<\infty$) (Joe, 1997, p.141-142).

Clayton copula density:

$$c_C(u_1,u_2,;\theta)=(1+\theta)(u_1,u_2)^{-1-\theta}(u_1^{-\theta}+u_2^{-\theta}-1)^{-\frac{1}{\theta}-2}, \tag{A.3}$$

for $0\le\theta<\infty$) and $\theta\ne0$.

Rotated versions of the Clayton, Gumbel, and Joe copulas are obtained by Equation A.4.

$$\begin{aligned} C_{90}(u_1,u_2) &= u_2 - C(1-u_1,u_2), \\ C_{180}(u_1,u_2) &= u_1 + u_2 - 1 + C(1-u_1,1-u_2), \\ C_{270}(u_1,u_2) &= u_1 - C(u_1,1-u_2). \end{aligned} \tag{A.4}$$

Figure A.1 displays contour plots for the normal, Student's t, Clayton, Gumbel, Frank and Joe copula with given dependency parameters.

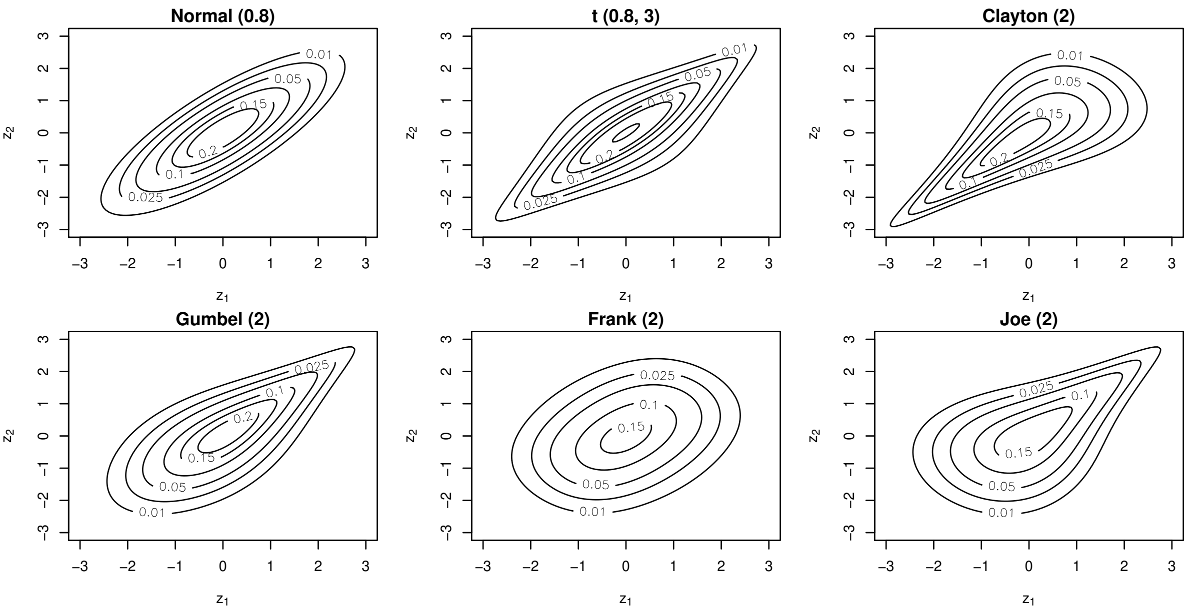

Figure A.1: Contour plots of density functions for some copula families.

A.3 Copula Estimation

The log-likelihood function for all parameters of the conditional copula distribution, i.e. of Equation 2.19:

$$log\big(f_t(\epsilon_{i,t},\epsilon_{j,t}|\Im_{t-1})\big) = log c_t(u_{i,t},u_{j,t}|\Im_{t-1},\theta_c) + log g_i, t(\epsilon_{i,t}|\Im_{t-1},\theta_i) log g_{j,t}(\epsilon_{j,t}|\Im_{t-1},\theta_j) \tag{A.5}$$

The estimator in the two-stage maximum likelihood framework, denoted $\hat{\theta} = [\hat{\theta}_i, \hat{\theta}_j, \hat{\theta}_c]$, assuming all samples run from $t = 1$ to $t = T$.

$$\begin{aligned} \hat{\theta}_i &= \arg\max_{\theta_i} \sum_{t=1}^{T} \log g_{i,t}(\epsilon_{i,t}|\Im_{t-1},\theta_i), \\ \hat{\theta}_j &= \arg\max_{\theta_j} \sum_{t=1}^{T} \log g_{j,t}(\epsilon_{j,t}|\Im_{t-1},\theta_j), \\ \hat{\theta}_c &= \arg\max_{\theta_j} \sum_{t=1}^{T} \log c_t(u_{1,t},u_{2,t}|\Im_{t-1},\hat{\theta}_i,\hat{\theta}_j,\theta_c). \end{aligned} \tag{A.6}$$

B | Estimation

B.1 Data

In Table B.1 information on the price series used in the analysis is given.

Price series	Name	Symbol	Contract size	Currency
SA S	Fish Pool Index Spot Salmon NOK/KG	FSPWKSP	NA	NOK
SA F	Fish Pool Salmon TRc1 NOK/KG	FSPFWDM	1.0 kilogram	NOK
SM S	Soyameal USA 48% Protein $/MT	SOYMUSA	NA	USD
SM F	ECBOT-Soybean Meal Continuous	CZMCS00	100.0 tons	USD
WH S	Wheat US HRS 14% Del Mineapolis/Dulut	WHTHRMD	NA	USD
WH F	MGE-WHEAT CONTINUOUS	MMWCS00	5000.0 bushels	USD
RO S	Rapeseed Oil EU Ex Mill FOB Rdam M	RPOLRDE	NA	EUR
RO F	Rapeseed Oil Dutch FOB NWE 1mth fwd	RPOLDNE	1.0 tonne	EUR

Table B.1: Metadata of price series obtained from Thomson Reuters Datastream.

Table B.2 presents descriptive statistics and Table B.3 gives tests performed on the log-returns series for the out-of-sample period.

Returns series	Mean	Median	Min.	Max.	St.dev	Skewn.	Exc.kur.
Salmon spot	–0.2433	–0.4511	–16.3797	17.2947	6.8191	0.2179	–0.4230
Salmon futures	–0.2743	–0.0638	–18.4922	13.1192	5.2660	–0.7051	2.7326
Soymeal spot	–0.3094	–0.4576	–6.3266	8.1932	2.4040	0.5296	1.2019
Soymeal futures	–0.3004	–0.5772	–6.3569	8.2326	2.3951	0.5081	1.1899
Wheat spot	–0.2704	–0.2605	–14.9733	13.7109	4.2348	–0.0391	1.6824
Wheat futures	–0.1872	–0.1722	–6.4513	6.8098	2.8215	–0.0169	–0.5228
Rapeseed oil spot	0.0543	0.0071	–9.0522	5.0208	2.0762	–0.9432	3.4196
Rapeseed oil futures	0.0526	0.0000	–9.9750	6.1036	2.3236	–0.7792	3.1113

Table B.2: Descriptive statistics for out-of-sample weekly spot and futures percentage log-returns.

Table B.3: Out-of-sample test statistics.

Returns series	**JB**	**ADF**[1]	**LBQ**[23]	**LM**[4]
Salmon spot	1.449	-2.55^{**}	13.63 (0.968)	23.0^{***}
Salmon futures	43.93^{***}	-2.25^{**}	25.94 (0.411)	70.8^{***}
Soymeal spot	12.15^{***}	-2.91^{***}	16.11 (0.912)	40.8^{***}
Soymeal futures	11.62^{***}	-3.02^{***}	18.39 (0.826)	42.3^{***}
Wheat spot	13.66^{***}	-2.85^{***}	25.88 (0.414)	18.6^{***}
Wheat futures	0.980	-3.98^{***}	23.40 (0.554)	16.0^{***}
Rapeseed oil spot	70.33^{***}	-2.10^{**}	22.11 (0.630)	63.1^{***}
Rapeseed oil futures	56.06^{***}	-1.77^{*}	28.12 (0.303)	82.9^{***}

Note: Tests applied are Jarque-Bera (JB), augmented Dickey-Fuller (ADF), Ljung-Box Q (LBQ) and Engle's Lagrange multiplier (LM) tests. ***, **, * denotes significance at the 1%, 5%, 10% level respectively.

Figure B.1 and B.2 show weekly log-returns for all commodity prices relevant in the analysis.

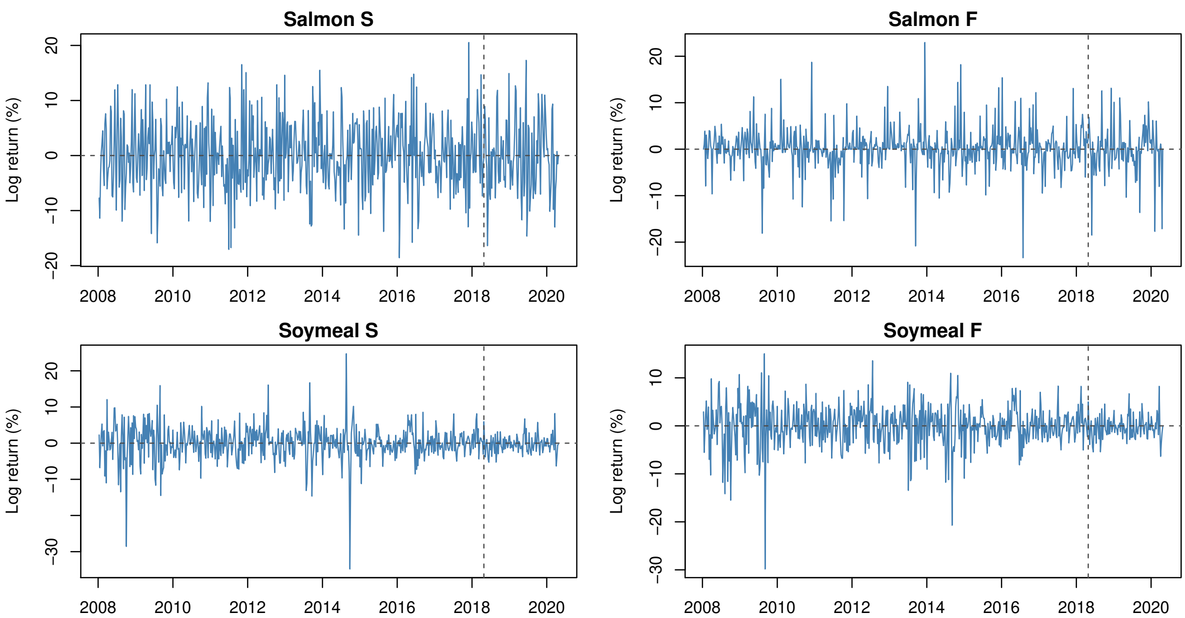

Figure B.1: Spot and futures log-returns for salmon and soymeal.

[1]Lag length $k = 12$ is chosen based on the commonly used rule of thumb by Schwert (2002), which is to choose $k = int\{12(T/100)^{1/4}\}$, where T denotes sample size.
[2]Values in parentheses are p-values.
[3]Lag length 25 is chosen. However, results are stable across a wide range of lag lengths.
[4]Results are for lag length 4.

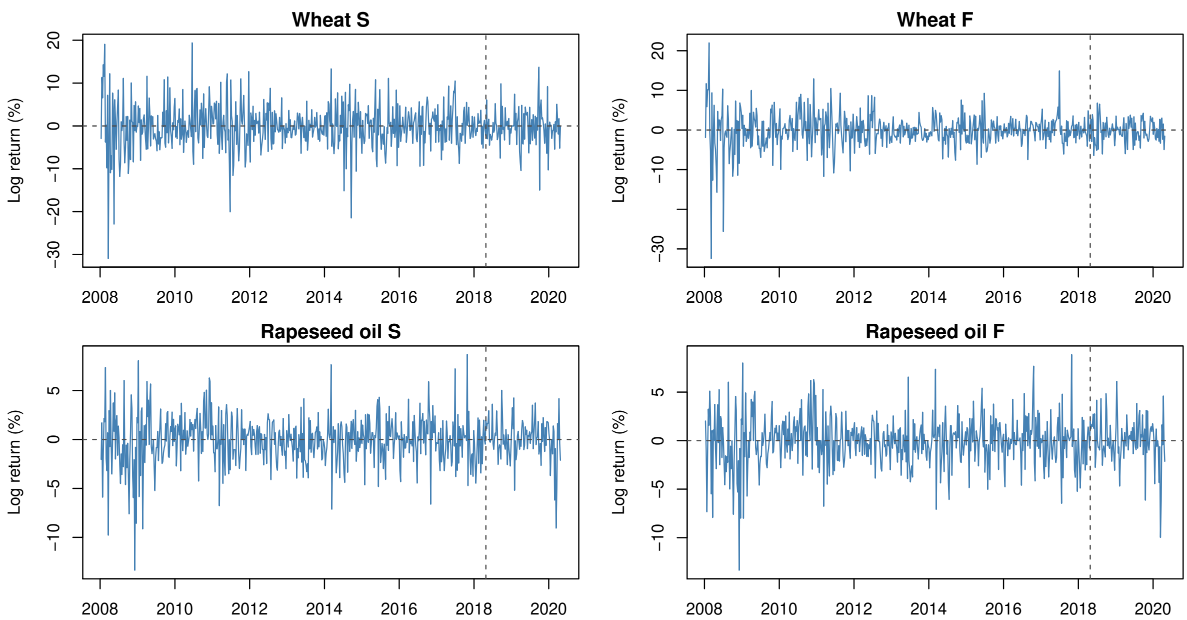

Figure B.2: Spot and futures log-returns for wheat and rapeseed oil.

Figure B.3 shows the ACF and PACF plots for salmon spot returns. The panels on the left show significant autocorrelation for lag 2 before fitting the AR model. The model seems to capture this, as seen in the panels on the right, and there is no significant autocorrelation left in the residuals of the AR models. We have equivalent results for all series. Figure B.4 shows the residuals of the AR model fitted to salmon spot returns.

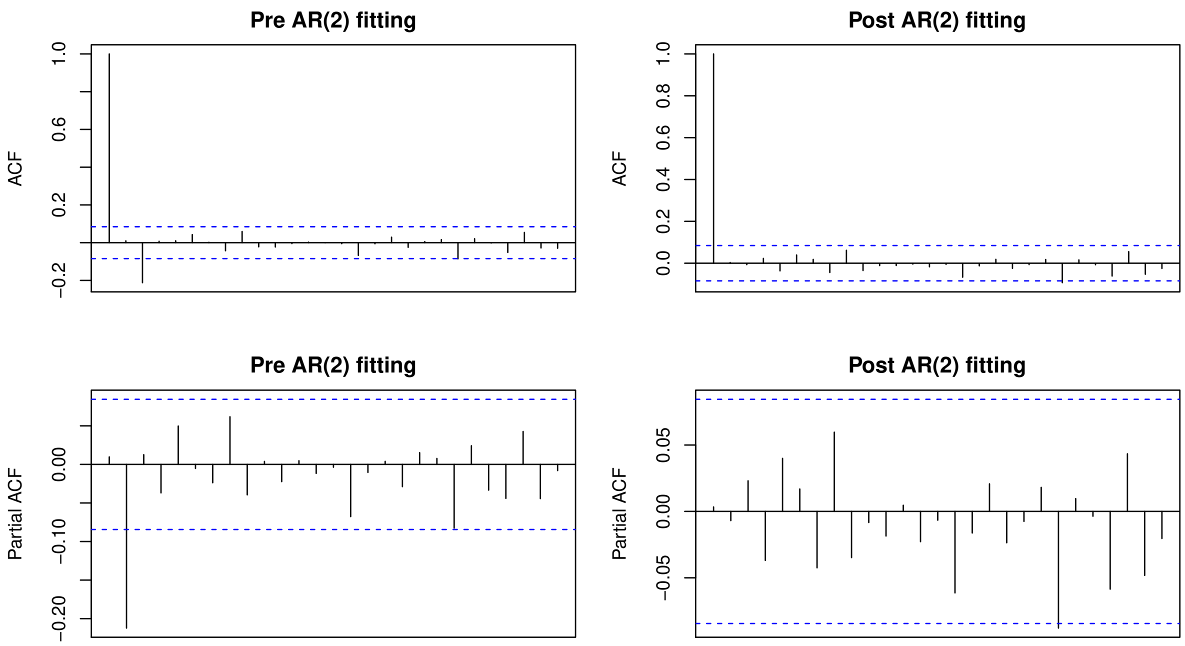

Figure B.3: ACF and PACF plots of salmon spot returns (left) and residuals of AR(2) model fitted to salmon spot returns (right).

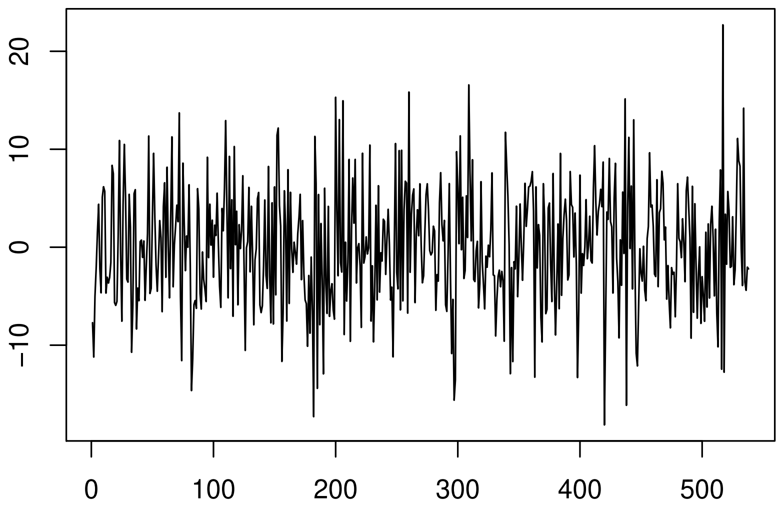

Figure B.4: Residuals of AR model fitted to salmon spot returns.

B.2 Estimated GARCH(1,1) models

The following figures show diagnostic plots for estimated GARCH(1,1) models in the form of ACF of squared standardised residuals, the empirical density and QQ plots.

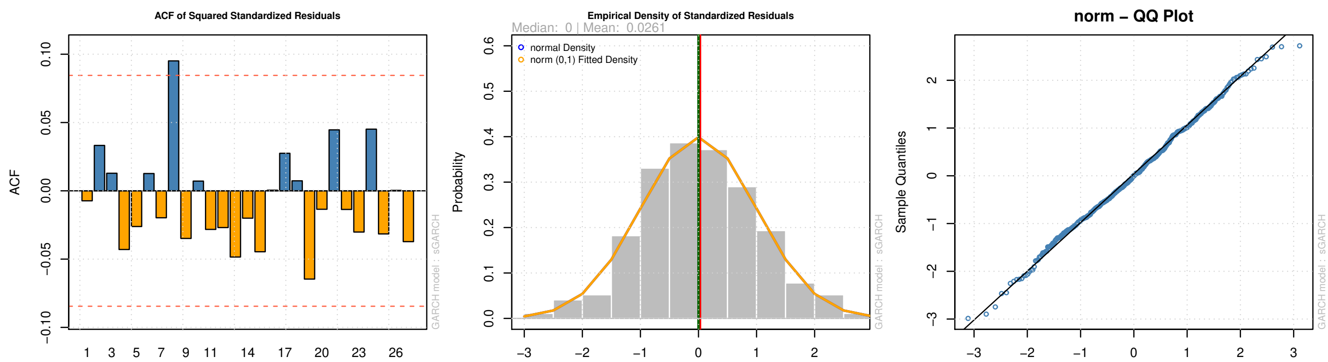

Figure B.5: GARCH(1,1) diagnostic plots, **salmon spot** series.

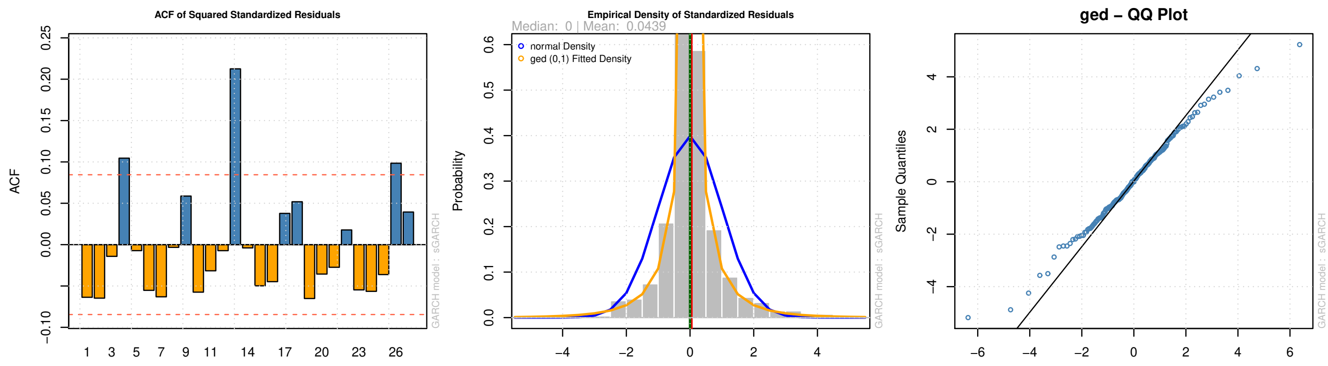

Figure B.6: GARCH(1,1) diagnostic plots, **salmon futures** series.

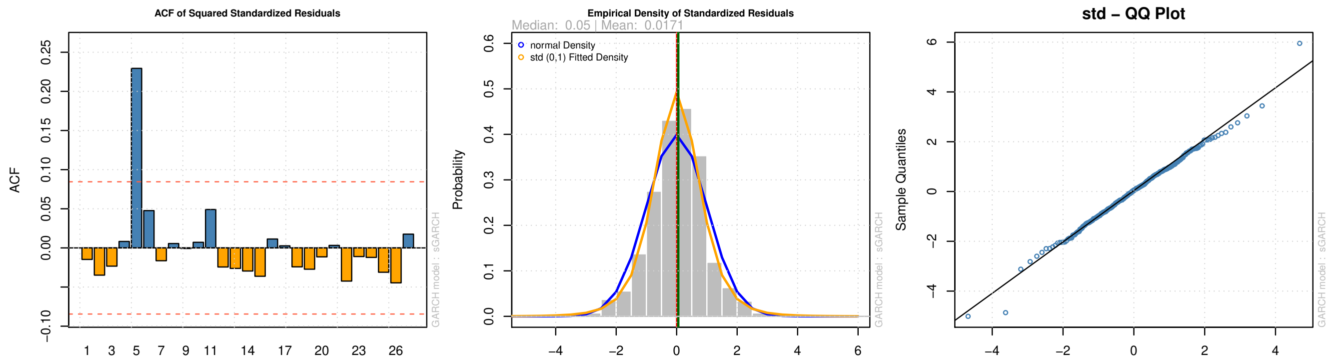

Figure B.7: GARCH(1,1) diagnostic plots, **soymeal spot** series..

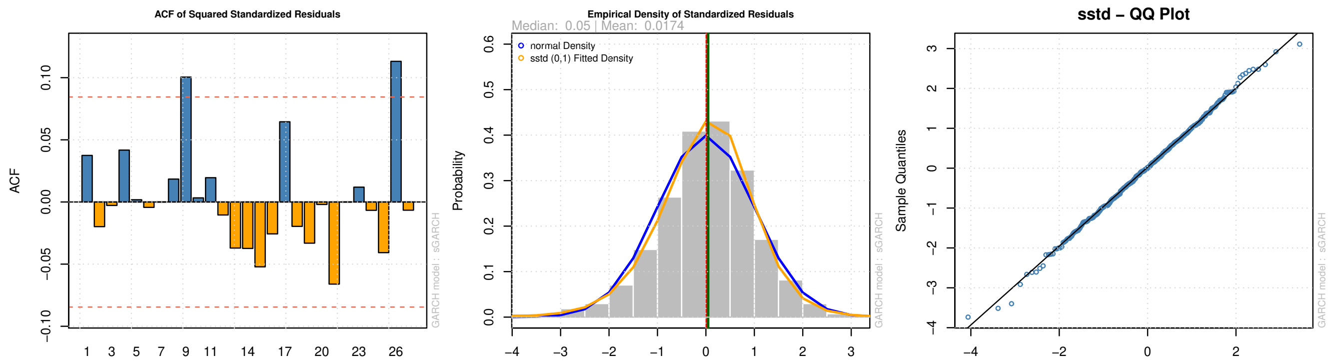

Figure B.8: GARCH(1,1) diagnostic plots, **soymeal futures** series.

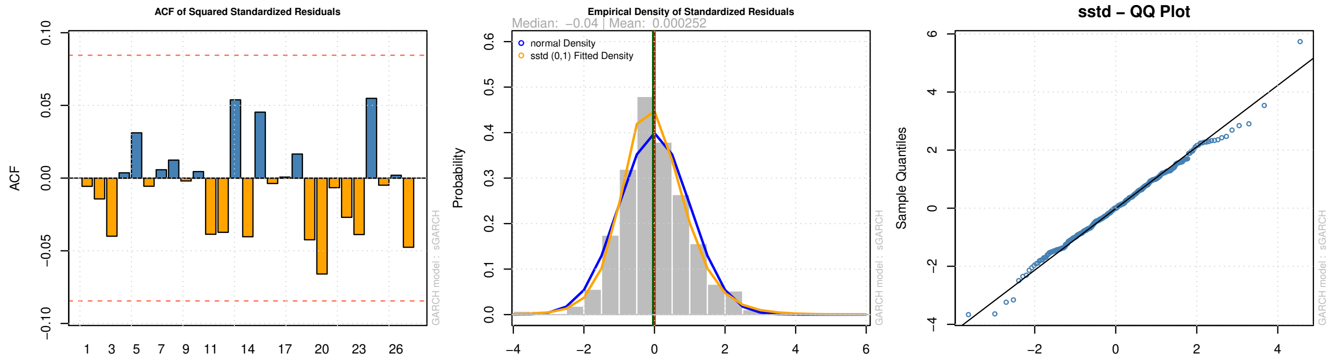

Figure B.9: GARCH(1,1) diagnostic plots, **wheat spot** series.

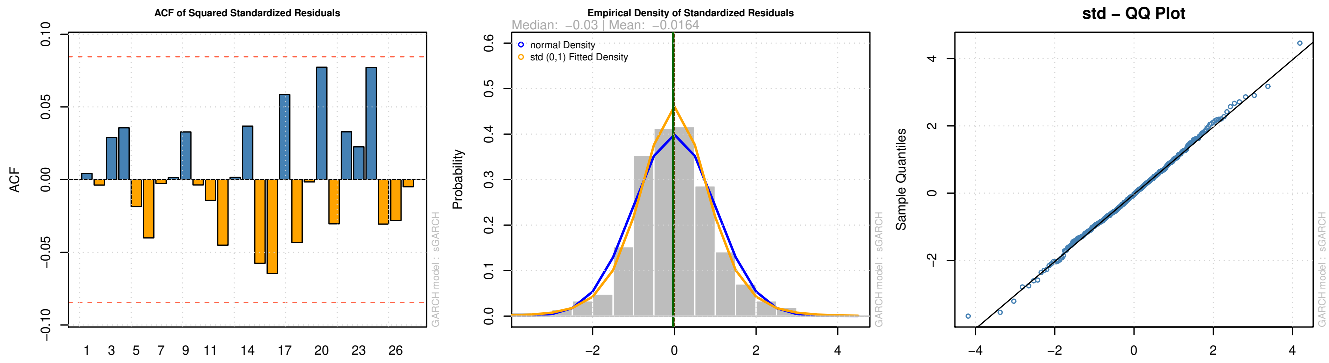

Figure B.10: GARCH(1,1) diagnostic plots, **wheat futures** series.

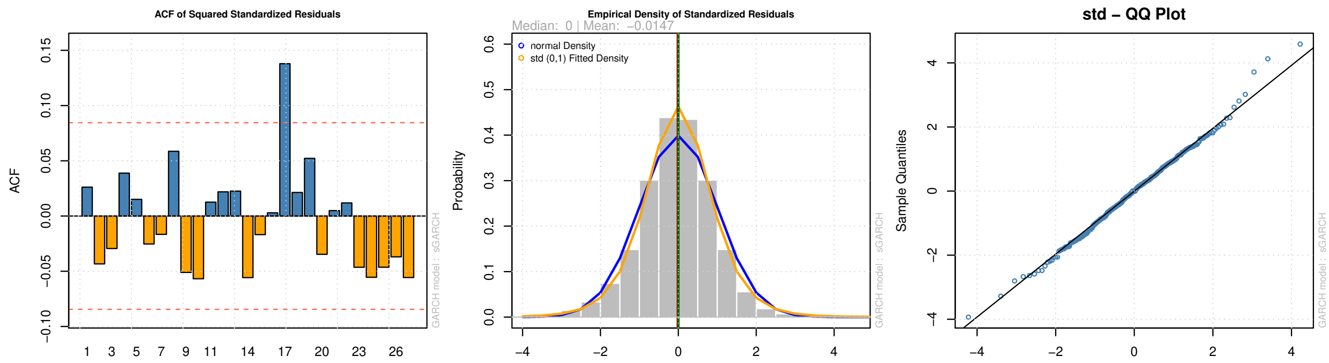

Figure B.11: GARCH(1,1) diagnostic plots, **rapeseed oil spot** series.

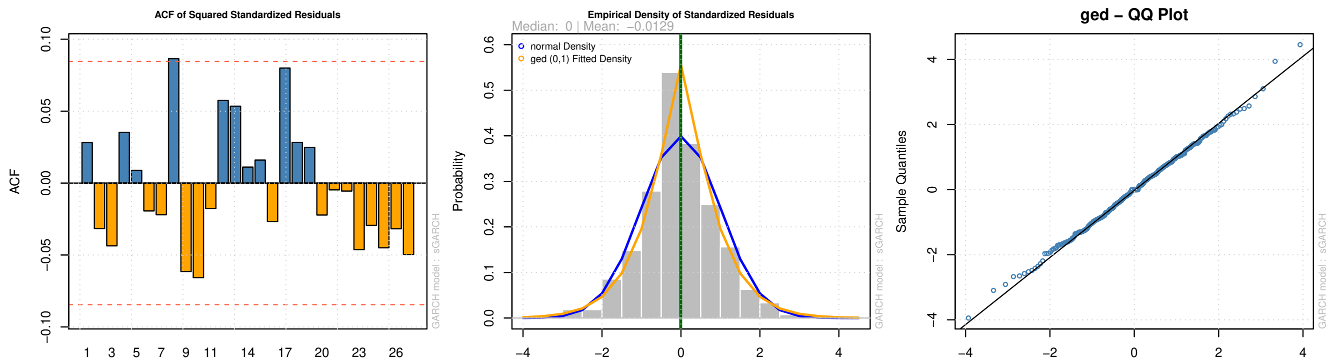

Figure B.12: GARCH(1,1) diagnostic plots, **rapeseed oil futures** series.

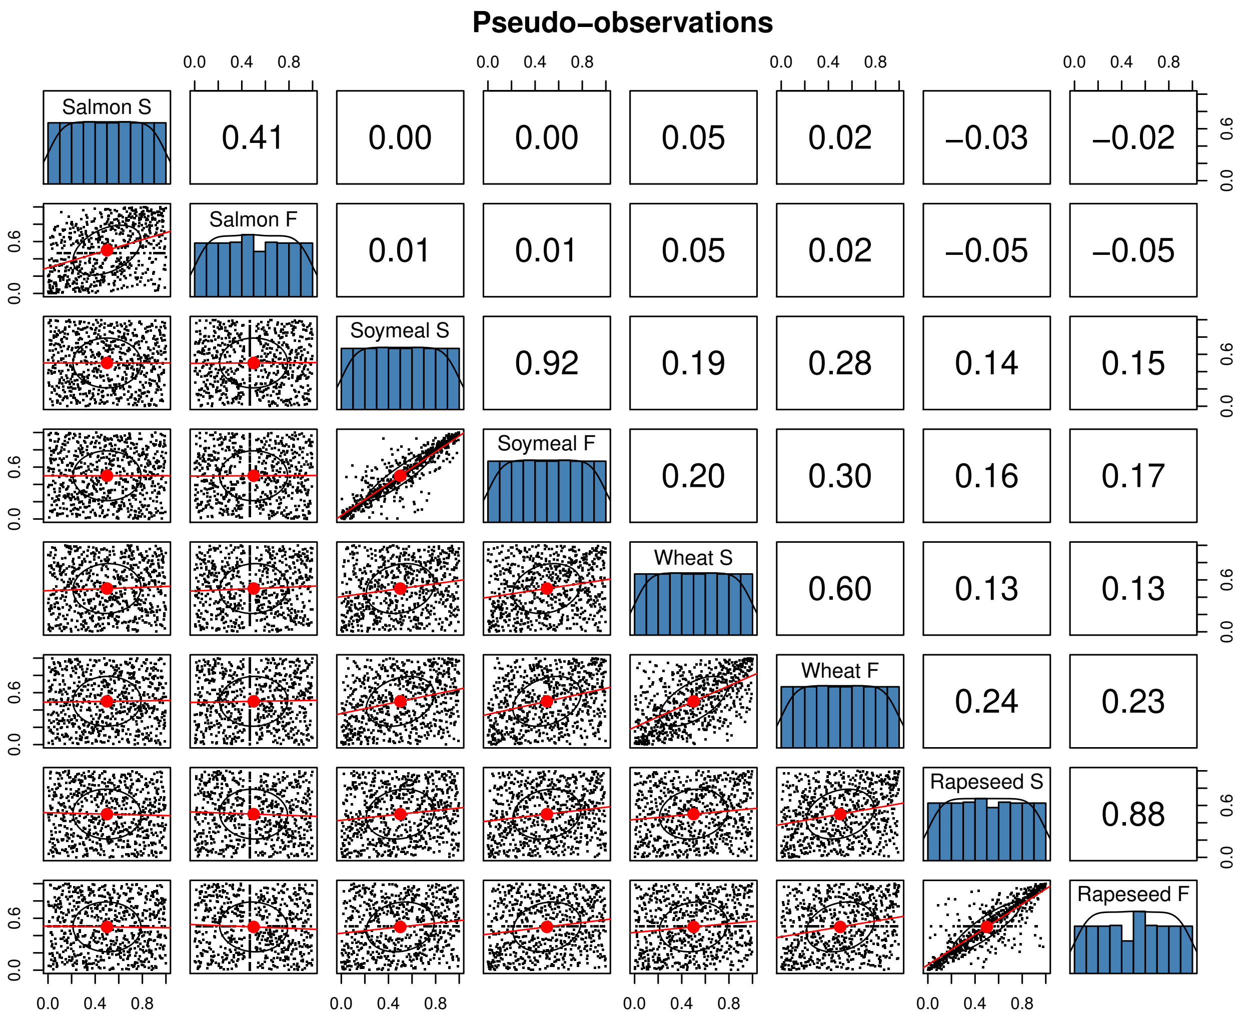

Figure B.13: Pseudo-observations obtained from standardised GARCH(1,1) residuals by the EDF, Equation 2.15.

C | Results

C.1 Estimated out-of-sample hedge ratios

Table C.1 shows out-of-sample mean, minimum, maximum and standard deviation (SD) for each model and commodity. For ease of comparison and consistency, all hedge ratios are presented from the same perspective, i.e. a *positive* ratio indicates a futures position *opposite* of the spot market, and a *negative* ratio indicates a futures position in the *same* direction as the spot market.

Table C.1: Statistics for estimated out-of-sample hedge ratios.

	DCC		**C-DCC**		**TVC**	**RWC**	
	Single	Multi	Single	Multi	Single	Single	Multi
Salmon							
Mean	0.584	0.571	0.559	0.544	0.522	0.664	0.588
SD	0.076	0.073	0.072	0.069	0.141	0.095	0.136
Min	0.454	0.447	0.430	0.419	0.254	0.436	0.274
Max	0.861	0.848	0.802	0.770	1.160	0.966	0.944
Soymeal							
Mean	0.951	1.078	0.977	1.141	1.000	1.044	-0.099
SD	0.033	0.251	0.034	0.224	0.083	0.039	0.911
Min	0.868	0.485	0.884	0.594	0.334	0.936	-2.027
Max	1.025	1.606	1.050	1.713	1.089	1.114	1.570
Wheat							
Mean	0.874	1.006	0.886	0.908	0.903	0.838	0.943
SD	0.223	0.532	0.227	0.531	0.243	0.203	2.190
Min	0.446	-0.152	0.446	-0.193	0.440	0.454	-2.034
Max	1.529	2.154	1.550	1.986	1.620	1.423	4.607
Rapeseed oil							
Mean	0.825	1.056	0.819	0.910	0.802	0.765	1.532
SD	0.067	0.490	0.086	0.525	0.147	0.092	2.011
Min	0.644	-0.421	0.597	-0.549	0.168	0.540	-2.381
Max	0.993	2.208	1.010	2.168	1.029	0.938	4.783

Note: Smallest and largest standard deviation (SD) for each commodity are marked red and blue respectively.

C.2 Various result plots

Figure C.1 shows the distributions of returns for different hedging strategies for the out-of-sample period. Single-hedges are shown in the left panel and multi-hedges in the right panel.

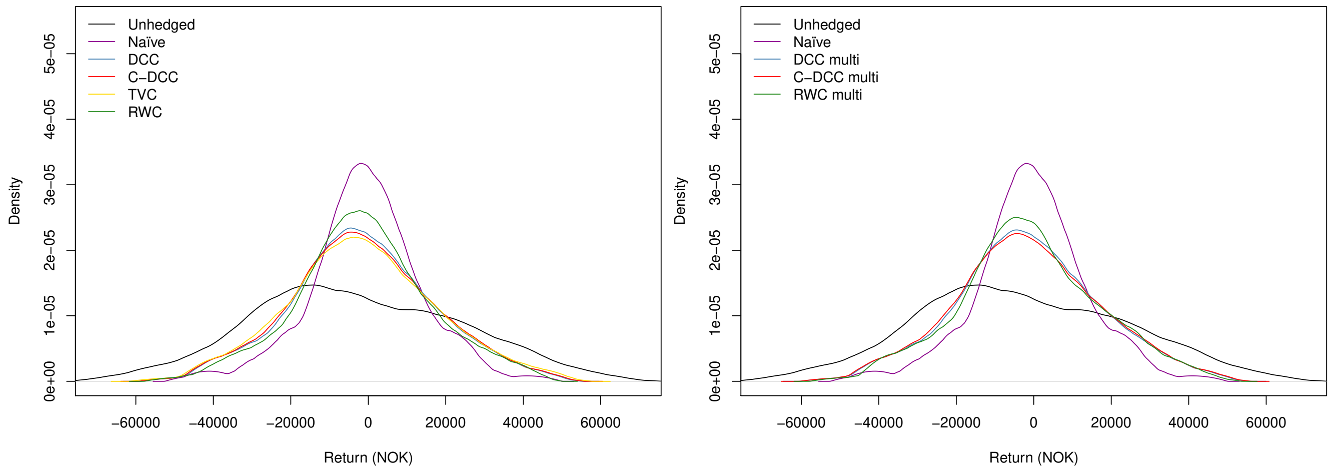

(a) Static and dynamic single-hedges. **(b)** Static and dynamic multi-hedges.

Figure C.1: Out-of-sample return distributions for the four-week hedging horizon.

C.3 Hedging paths

Figures C.2 to C.4 show the realised return paths. The dashed vertical line separates the in- and out-of-sample period.

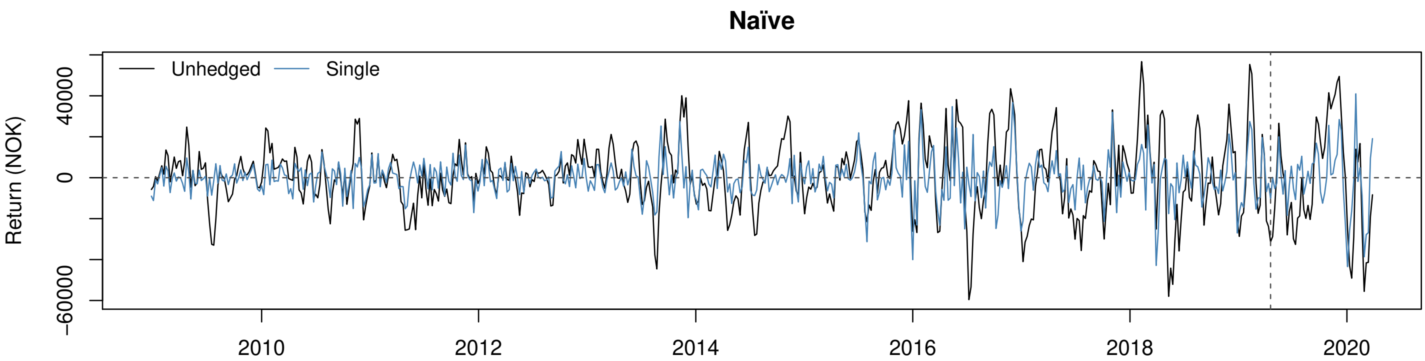

Figure C.2: Hedging path of the naïve hedge for the four-week horizon.

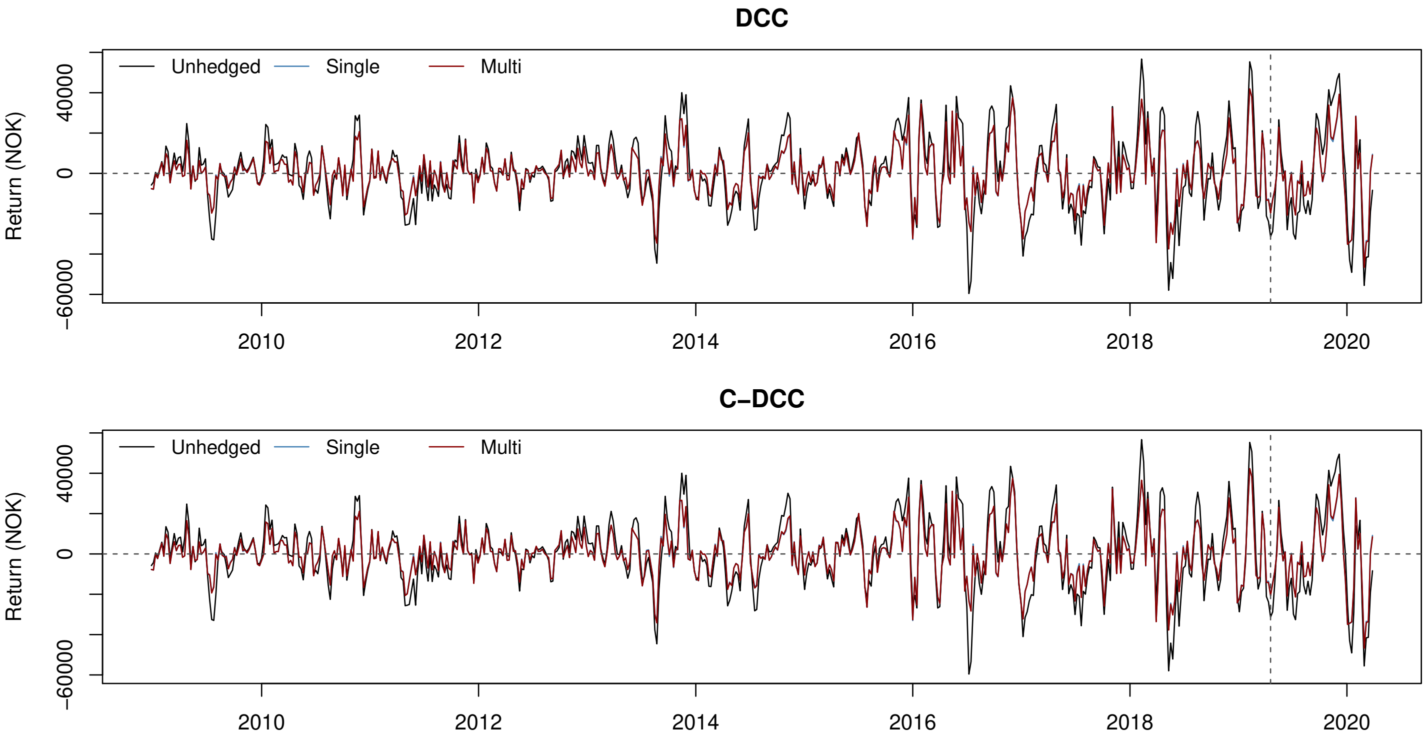

Figure C.3: Hedging path of the DCC and C-DCC models for the four-week horizon.

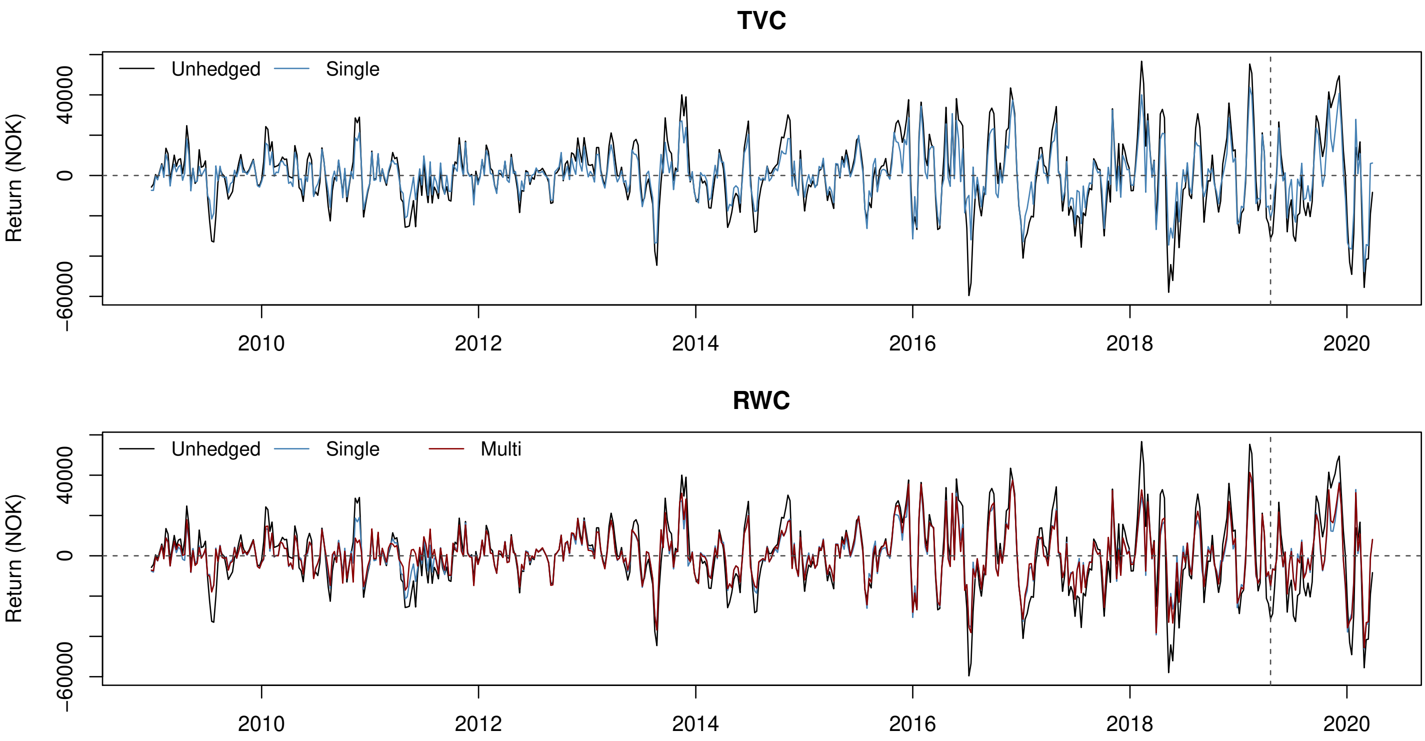

Figure C.4: Hedging path of the TVC and RWC models for the four-week horizon.

C.4 Software setup

Software: RStudio v.1.2.5033[1] (R version 3.6.3).

External R Packages used:

- abind v.1.4-5 (Plate and Heiberger, 2016).
- aTSA v.3.1.2 (Qiu, 2015).
- copula v.1.0-0 (Hofert et al., 2020).
- cubature v.2.0.4 (Narasimhan et al., 2019).
- doParallel v.1.0.15 (Plate and Heiberger, 2016).
- fgarch v.3042.83.2 (Wuertz et al., 2020).
- foreach v.1.5.0 (Ooi, Microsoft and Weston, 2019).
- forecast v.8.12 (Hyndman et al., 2020).
- psych v.1.9.12 (Revelle, 2019).
- quantmod v.0.4.17 (Ryan et al., 2020).
- RColorBrewer v.1.1-2 (Neuwirth, 2014).
- readxl v.1.3.1 (Wickham et al., 2019).
- rmgarch v.1.3-7 (Ghalanos, 2019a).
- rugarch v.1.4-2 (Ghalanos, 2020).
- skewt v.0.1 (King, 2012).
- tseries v.0.10-47 (Trapletti and Hornik, 2019).
- VineCopula v.2.3.0 (Nagler et al., 2019).

Note:
Code used to produce the estimates and results in the thesis are given upon request.

[1] "Orange Blossom" (330255dd, 2019-12-04), Mozilla5.0 (Windows NT 10.0 Win64 x64) AppleWebKit537.36 (KHTML, like Gecko) QtWebEngine5.12.1 Chrome69.0.3497.128 Safari537.36

NTNU
Norwegian University of
Science and Technology