ARTICLE IN PRESS

Journal of Commodity Markets xxx (xxxx) xxx

Contents lists available at ScienceDirect

Journal of Commodity Markets

journal homepage: www.elsevier.com/locate/jcomm

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

Aleksander H. Haarstad, Maria Lavrutich *, Kristian Strypet, Eivind Strøm

Norwegian University of Science and Technology, 7491, Trondheim, Norway

ARTICLE INFO

JEL classification:
C32
G17
G32
Q22

Keywords:
Aquaculture
Salmon farming
Salmon feed
Risk management
Multi-commodity hedging
Copulas

ABSTRACT

Cost management has received limited attention in the aquaculture industry due to historically high profit margins. This trend, however, is not likely to continue. This creates a need for knowledge on optimally managing financial risks. In this study, we address the joint input-output price hedging problem of salmon farmers. Along with salmon, we consider three essential commodities used in fish feed mixtures. We use state-of-the-art copula models to examine multi-commodity hedging strategies. Our results show significant potential in reducing the joint price risk. Our key finding is that multi-commodity hedging improves hedging effectiveness for short horizons and risk-return trade-off for longer horizons. Salmon farmers face a trade-off where longer hedging horizons yield increased effectiveness and lower costs, yet require increased pre-planning of slaughtering volumes.

1. Introduction

Improvement of risk management practices is a crucial step towards achieving economic sustainability and profitability of the salmon farming industry. 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 (Asche et al., 2008). The growth potential of the industry has been, however, limited by biological factors. The biological nature of the salmon farming business leads to periods of higher mortality rates and periods of forced excessive slaughtering (Abolofia et al., 2017; Pincinato et al., 2021). This, together with seasonality in growth and harvesting, result in large variations in salmon supply which carry over into financial markets and contribute to volatile prices (Oglend, 2013; Thyholdt, 2014; Asche, 2017). Another contributing factor to increased volatility are spillover effects from other markets (Dahl and Jonsson, 2018a,b; Dahl and Yahya, 2019). This volatility has been further exacerbated by the COVID-19 outbreak. The salmon prices dropped 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.

The largest players in the Norwegian market, for example, 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, 2020). Such contracts can serve as means for risk transfer from those who wish to reduce risk, typically

* Corresponding author.
E-mail address: (A.H. Haarstad), maria.lavrutich@ntnu.no (M. Lavrutich), (K. Strypet), (E. Strøm).

https://doi.org/10.1016/j.jcomm.2021.100182
Received 12 October 2020; Received in revised form 15 February 2021; Accepted 24 February 2021
Available online XXX

a salmon farmer, to those with a higher risk appetite (Asche et al., 2016a). The study by Schütz and Westgaard (2018) indicates that farmers should use futures contracts at fairly low degrees of risk-aversion. The salmon price has also been studied in the price forecasting literature, most recently by Steen and Jacobsen (2020). In addition to the salmon price, however, optimizing business performance requires successful management of costs and related risks. The main input cost for salmon producers is fish feed (Iversen et al., 2020). 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. Today, feed producers are offering feed price hedging services for buyers. 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.[1] Nonetheless, findings in several studies, e.g. Smith and Stulz (1985); Graham and Smith (1999), suggest that reducing exposure can add significant value. This creates 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 standardized financial hedging tools such as futures on the feed itself and limited knowledge of the potential and use of financial hedging among industry players.

In this paper, 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 providing practical steps towards better risk management practices in the industry by applying advanced techniques to a stylized scenario applicable across different value chain set-ups. We move 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 fish feed 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 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.

Our second contribution is extending the current literature on hedging salmon price risk by applying a multivariate GARCH model to obtain dynamic hedge ratios for both salmon and fish feed commodities. Additionally, we analyze 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 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. Dahl and Oglend (2014) showed that volatility for fish feed ingredients such as fish meal is higher than the salmon price. However, 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 (Veikkaus Oy, 2015). On the other hand, successful cross-hedging of fish meal suggests similar approaches should be examined for fish feed hedging, which is the focus of our paper.

Lastly, we contribute to the aquaculture literature by investigating the potential of state-of-the-art multi-commodity hedging methods in salmon markets. The study is an extension of contributions to output price hedging such as Misund and Asche (2016) and Schütz and Westgaard (2018). Multi-commodity hedging has to the best of our knowledge not yet been studied in the context of salmon farming. This paper fills the gap related to modeling of input hedging, as well as joint input and output hedging in the current aquaculture risk management literature.

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 minimizing 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 focus on multiple commodities, we overcome these issues by using multiple parametric copulas for describing the dependence between two variables, which in many cases can be more effective than linear

[1] This insight was shared by an actor in the salmon farming industry. Furthermore, annual reports from the actors show no evidence of hedging in the feed input commodity markets.

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 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. 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 modeling hedge ratios in the aquaculture industry, which is confirmed by our study. We show that hedging outcomes are significantly improved 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. Vice versa, marginal distributions can be combined with a copula function to form a joint multivariate distribution, which we utilize 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 (Jondeau and Rockinger, 2006; Power and Vedenov, 2008). Copulas can therefore provide realistic joint distributions, which can be exploited in risk management by obtaining more realistic GARCH models. Furthermore, dynamic copula models yield greater flexibility in capturing vital features in dependence structures, increasing the robustness of the results (Aepli et al., 2017). 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. Sections 2 and 3 present the applied methodology. A description of our application to the salmon industry, data and estimated models are presented in Section 4. Results are discussed in Section 5. Section 6 concludes the paper.

2. Methodology

A widely used technique for managing price risk is hedging with futures contracts. Consider a 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 minimize 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 $h = 1$. This strategy is based on the assumptions 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). An alternative to the naïve hedge is to find the *optimal hedge ratio*, h^*, which minimizes 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*.[2] Given that Asche et al. (2016b) finds 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 paper 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

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

Following Brooks (2014, p.465–466), we derive the variance minimizing dynamic hedge ratio which is given by

$$h_t^* = \frac{Cov_t(s_t, f_t)}{Var_t(f_t)}, \tag{2}$$

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. Then in order to find the optimal time-varying hedge ratios, one needs to estimate the conditional-variances and covariances for spot and futures price returns in the portfolio.

While the hedge ratio in (2) is optimal when considering the spot and futures price returns of a single commodity, this is not

[2] 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 et al. (2016b).

necessarily the case for a multi-commodity setting 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 minimizing portfolios of each commodity, hedged independently with hedge ratios as given by (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 prevents speculative positions when the spot and futures markets are positively correlated.[3] We denote the vector of optimal dynamic hedge ratios when considering i commodities hedged separately as $\mathbf{h}_{S,t} = \{h_{1,t}, \dots, h_{i,t}\}$, 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 unfavorable 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 $\mathbf{h}_{M,t} = \{h_{1,t}, \dots, h_{i,t}\}$ and will henceforth refer to it as the multi-hedge ratio, given by[4]

$$\mathbf{h}^*_{\mathbf{M,t}} = \left[diag(\mathbf{Q})\right]^{-1} \left[\sum_{FF}(t)\right]^{-1} \sum_{FP}(t)\mathbf{Q}, \tag{3}$$

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 and $diag(\mathbf{Q})$ is a diagonal matrix with $\mathbf{Q}$ on the main diagonal (Fackler and McNew, 1993).

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. The 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 when 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 minimize the variance of returns, HE is measured as the percentage reduction of variance in the hedged portfolio against the unhedged portfolio (Ederington, 1979), given by

$$\text{Hedge effectiveness} = 1 - \frac{Var(\text{Hedged portfolio})}{Var(\text{Unhedged portfolio})}. \tag{4}$$

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 (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):

$$ES^{\alpha}(X) = \left(-\frac{1}{\alpha}\right)\left(\mathbb{E}\left[X\mathbb{I}_{X \le x^{\alpha}}\right] - x^{\alpha}\left(\mathbb{P}\left[X \le x^{\alpha}\right] - \alpha\right)\right). \tag{5}$$

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

$$TCE^{\alpha}(X) = -\mathbb{E}\{X|X \le x^{\alpha}\}. \tag{6}$$

3. Models

3.1. GARCH models

In order to obtain time-varying hedge ratios and capture important characteristics such as heteroscedasticity we estimate GARCH models. Consider a time series of commodity prices with a sample of T observations. Let r_t denote the continuously compounded return, or log-change, between prices at time t and $t-1$. According to the GARCH model, the return at time t can be expressed as

$$r_t = \mu_t + \sigma_t \epsilon_t, \quad where \quad \epsilon_t \sim g(0, 1, \theta), \tag{7}$$

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

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

where μ_t is the conditional mean, σ_t^2 is the conditional variance,[5] and ϵ_t is the standardized residual at time t. Furthermore, $g(0, 1, \theta)$ is the assumed conditional distribution with distributional parameters θ. While ϵ_t typically is assumed to be standard normal, we also consider the generalized error distribution (GED), the Student's t distribution and the skewed t distribution proposed by Fernández and Steel (1998).

The conditional variance of (7) is modeled as

$$\sigma_t^2 = \omega + \beta h_{t-1}^2 + \alpha u_{t-1}^2, \tag{8}$$

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 $\mathrm{Var}(u_t) = \frac{\omega}{1-(\alpha+\beta)}$, where we require $(\alpha + \beta) < 1$ to ensure stationarity. While the GARCH(1,1) model can be extended to a GARCH(p, q) model, the (1,1)-specification is generally sufficient when an appropriate distribution for the residuals is specified (Brooks, 2014; Hansen and Lunde, 2005).

In order to capture time-varying correlations between commodities, we use the dynamic conditional correlation (DCC) model proposed by Engle (2002) and Tse and Tsui (2002) as a baseline. The DCC model is not subject to the curse of dimensionality, and thus allows us to use high dimensional data. This is relevant for commodity processors, such as salmon farmers, hedging against the risk of multiple commodities.[6] For this model 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 $\mathbf{H}_t$ is defined as

$$\mathbf{H}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t, \tag{9}$$

where $\mathbf{D}_t$ is a diagonal matrix containing the conditional standard deviations obtained from (8) for each individual series, and $\mathbf{R}_t$ is the conditional correlation matrix. Both $\mathbf{D}_t$ and $\mathbf{R}_t$ vary over time, producing a new variance-covariance matrix for each time step, differentiating DCC from the constant conditional correlation (CCC) model.

One of the shortcomings of DCC models is the requirement in the maximum likelihood estimation (MLE) procedure that the standardized residuals follow the multivariate normal distribution. This may not be consistent with financial data, which often contain features such as skewness and excess kurtosis. To address this issue, we apply a copula approach (Patton, 2006b).

3.2. 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 $X = (x_1, \ldots, x_n)$ with margins $G_1(x_1), \ldots, G_n(x_n)$, can be decomposed as $F(x_1, \ldots, x_n) = C(G_1(x_1), \ldots, G_n(x_n))$.

In this paper, we use parametric copula families from two main categories: elliptical copulas with symmetric dependency structures and 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. We provide further details on the respective copulas in Appendix A.1.

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-dependent data if the latter satisfy the Markov Property. 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 standardized GARCH residuals $\epsilon_{i,t}$, $\epsilon_{j,t}$ at time t can be written as

$$\begin{aligned} f_t(\epsilon_{i,t}, \epsilon_{j,t} | \Im_{t-1}) &= c_t\left(G_{i,t}(\epsilon_{i,t}|\Im_{t-1}), G_{j,t}(\epsilon_{j,t}|\Im_{t-1}) | \Im_{t-1}\right) \times g_{i,t}(\epsilon_{i,t}|\Im_{t-1}) \times g_{j,t}(\epsilon_{j,t}|\Im_{t-1}), \\ \text{where} \quad c_t(u_{i,t}, u_{j,t}|\Im_{t-1}) &= \frac{\partial^2 C_t\left(u_{i,t}, u_{j,t}|\Im_{t-1}\right)}{\partial u_{i,t} \partial u_{j,t}}, \end{aligned} \tag{10}$$

$g_{i,t}(\epsilon_{i,t}|\Im_{t-1})$ 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 (10) are estimated by a two-stage maximum likelihood procedure, in which parameters for the density functions of $\epsilon_{i,t}$, $\epsilon_{j,t}$ and parameters for the copula function are estimated (Patton, 2012). 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, $\sigma_{ij,t}^2$, can be generated by numerical integration. The conditional variances (obtained from the GARCH models) and the conditional covariances are used to compute the optimal hedge ratios by (2) and (3). In the next section we present the copula models we use to obtain the conditional joint density of (10).

[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.

[6] For example, the general BEKK model in the case of four commodities (8 time series), would require jointly estimating 164 parameters using the R package by Schmidbauer et al. (2016).

3.3. Copula GARCH models

In the following section we present three copula GARCH models. As opposed to GARCH models, copula based methods allow for more flexible modeling 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.

First, we consider an extension of the DCC model, namely the copula-DCC model (C-DCC), as described in Ghalanos (2019). 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 other words, we relax the assumption that the distribution of the conditional marginals are standard normal and allow a non-normal dependency structure.

Assume we have $i = 1, \dots, n$ conditional marginal distributions estimated in the first stage by GARCH(1,1) processes, where G_i is the conditional CDF of the ith 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

$$c_t(u_{1,t}, \dots, u_{n,t}|\mathbf{R}_t, \eta) = \frac{f_t\left(G_1^{-1}(u_{1,t}|\eta), \dots, G_n^{-1}(u_{n,t}|\eta)|\mathbf{R}_t, \eta\right)}{\prod_{i=1}^{n} f_i\left(G_i^{-1}(u_{i,t}|\eta)|\eta\right)}, \tag{11}$$

where $u_{it} = G_{i,t}(x)$ is the probability integral transform (PIT) of each series by its empirical distribution function (EDF), $G_i^{-1}(u_{i,t}|\eta)$ 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. The joint density of the two-stage estimation is given by

$$f_t(\mathbf{r}_t|\mathbf{h}_t, \mathbf{R}_t, \eta) = c_t(u_{i,t}, \dots, u_{n,t}|\mathbf{R}_t, \eta) \prod_{i=1}^{n} \frac{1}{\sigma_{i,t}} g_{i,t}(\epsilon_{i,t}|\theta_i), \tag{12}$$

where $\epsilon_{i,t} \sim g_i(0, 1, \theta_i)$ are the standardized 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, 2019).

Second, we consider time-varying copula models (TVC), initially proposed by Patton (2006a,b). In these models, 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. The time-varying parameter for the normal, Student's t and Gumbel copula are modeled as

$$\begin{aligned}
\text{Normal:} \quad & \rho_t = \Lambda\left(\omega_N + \beta_N \rho_{t-1} + \alpha_N \frac{1}{n} \sum_{k=1}^{n} \Phi^{-1}(u_{i,t-k})\Phi^{-1}(u_{j,t-k})\right), \\
\text{Student's t:} \quad & \rho_t = \Lambda\left(\omega_T + \beta_T \rho_{t-1} + \alpha_T \frac{1}{n} \sum_{k=1}^{n} T^{-1}(u_{i,t-k}; \nu)T^{-1}(u_{j,t-k}; \nu)\right), \\
\text{Gumbel:} \quad & \theta_t = \kappa\left(\omega_G + \beta_G \theta_{t-1} + \alpha_G \frac{1}{n} \sum_{k=1}^{n} |u_{i,t-k} - u_{j,t-k}|\right),
\end{aligned} \tag{13}$$

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 parameters in (13) are estimated by maximum likelihood.[7] The model, therefore, requires estimating (13) 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 (10).

Lastly, we consider 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 assume that the copula family modeling the distribution is constant.[8] An additional disadvantage of the DCC and C-DCC models is that in the multi-dimensional setting, they assume the same parametric copula for any pair of random variables in the model. To relax these assumptions, we consider each variable pair individually and select the best fitting conditional bivariate copula c_t among {normal, Student's t, Clayton, Gumbel, Frank, Joe}.[9] Therefore, we allow different dependency structures between different variable pairs that additionally vary through time.

The model estimation is done in two stages. The first stage consists of estimating the conditional marginal distributions by GARCH(1,1) models. The second stage consists of estimating conditional copulas c_t between variable pairs, which ultimately are used to obtain the conditional covariances of (2) and (3). We apply the estimation procedure to a moving window of N observations. I.e. $t = \{1, \dots, N\}$ observations are used to estimate densities at $t = N$, $t = \{2, \dots, N+1\}$ 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

[7] See Patton (2006a,b) for more details.

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

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

$\sigma^2_{ij,t}$, we estimate the bivariate parametric copula of each family by maximum likelihood. Pseudo-observations $u_{i,t}$, $u_{j,t}$ are obtained by the EDF using the standardized residuals estimated in the first stage. The best fitting copula family is then selected by the AIC criterion.

4. Estimation

We begin by constructing a conceptual hedging framework tailored to the salmon farming industry, taking into account both input and output commodity prices. We consider a hypothetical well-established Norwegian salmon farming company which has an objective of reducing the price exposures of its operations. 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. Second, we assume the company has biomass assets in all stages of the salmon production cycle, and an average volume of 3000 tons is slaughtered and sold every week. The same price is realized for all salmon sold within the same week. This also implies constant feed consumption throughout the year. As the farmers buy feed on a regular basis, we assume that price movements in the feed input commodity markets affect feed prices the same week.

Third, the company aims to reduce the exposure to price risk of both fish feed and salmon. Given that futures contracts on fish feed 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 sell feed on contracts that transfer the full risk exposure to the purchasing party, such as cost-plus type contracts (Mowi, 2019; BioMar Group, 2019). Even though some salmon production companies control fish feed costs by upstream integration, they are still exposed to the underlying input commodity price, making our hedging framework applicable across multiple value chain set-ups.

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 can vary to some degree (Aas et al., 2019).[10] A feed conversion ratio of 1.1, implies that 0.22 kg of soymeal, wheat and rapeseed oil (0.66 kg in total) are required for each 1 kg of salmon produced (Mowi, 2020). 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 5000.0 bu,[11] respectively (Parcell and Pierce, 2011). For the purpose of our analysis, we assume the contract sizes are sufficient for the company's hedging requirements.

In order to reduce price risk, it is preferable to enter a futures contract that is highly correlated to the spot price. 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, which is also the most liquid. Therefore, we use the front month contract for all commodities in our analysis. This is consistent with related studies such as Misund and Asche (2016); Bloznelis (2018). Several studies, however, find that salmon futures have been found to pose liquidity issues in recent years, most notably due to lower participation on the exchange (Oglend and Straume, 2020). For the sake of tractability, in this paper we disregard this problem and assume that salmon farmers enter and exit all futures positions at the same time in order to focus on comparisons among different hedging models. For further details on the futures liquidity and its effect on the hedging strategies see, e.g., Bloznelis (2018) and Andersen (2019).

In addition, we address the sensitivity of our results with respect to contract lengths by exploring different hedging horizons. Hedging effectiveness typically increases for co-integrated 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 various stochastic factors, which make them difficult to predict in the far future (Hagspiel et al., 2018; Brakstad et al., 2019). Therefore, some degree of production risk is unavoidable. Nonetheless, its effect on hedging results can be reduced by selecting an appropriate hedging horizon in which the salmon farmer is able to plan production volumes. The farmer can also adjust hedging positions in the case of extreme events. In this study, we select a four-week hedging duration as baseline, which is consistent with studies such as Misund and Asche (2016).

Lastly, we assume a fixed transaction fee including trading and clearing when calculating transaction costs for all commodity futures contracts (Fish Pool, 2020). 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.[12] This will to some extent compensate for extra costs associated with establishing memberships and licenses on international exchanges. With these underlying assumptions, we define the hedged portfolio return of input and output commodities as

[10] Fish meal and fish oil are left out as there are no futures contracts available. Additionally, the trend is showing a decreasing share of both of them in modern feed compositions (Aas et al., 2019).

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

[12] 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).

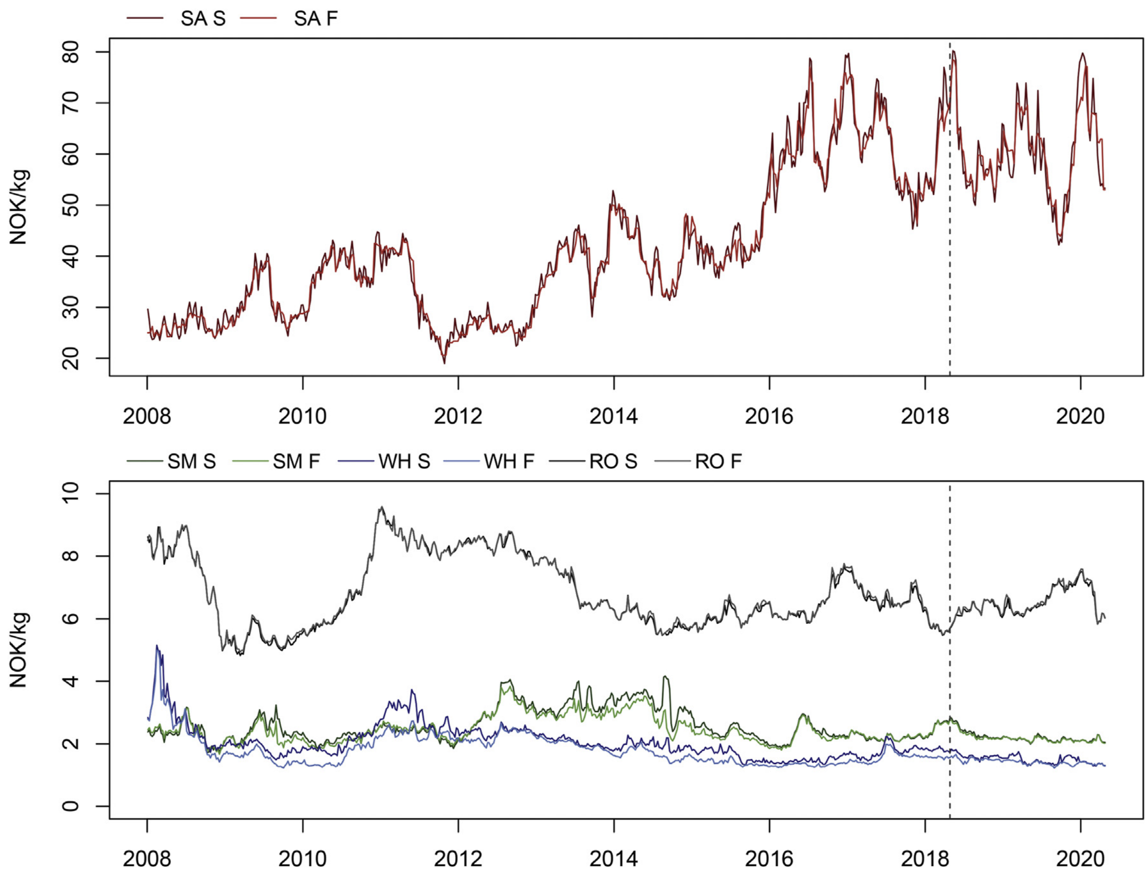

Fig. 1. Weekly spot and futures prices for salmon (SA), soymeal (SM), wheat (WH) and rapeseed oil (RO). S and F denotes spot and futures prices.

$$\begin{aligned}\pi(\mathbf{h}) = {} & Q^{SA}\left(s_1^{SA} - s_0^{SA}\right) - h^{SA}Q^{SA}\left(f_1^{SA} - f_0^{SA}\right) \\ & - Q^{SM}\left(s_1^{SM} - s_0^{SM}\right) + h^{SM}Q^{SM}\left(f_1^{SM} - f_0^{SM}\right) \\ & - Q^{WH}\left(s_1^{WH} - s_0^{WH}\right) + h^{WH}Q^{WH}\left(f_1^{WH} - f_0^{WH}\right) \\ & - Q^{RS}\left(s_1^{RO} - f_0^{RO}\right) + h^{RO}Q^{RO}\left(f_1^{RO} - f_0^{RO}\right),\end{aligned} \tag{14}$$

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 realized 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.

4.1. Data

In this section, we describe 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 and use the front month salmon futures prices. Weekly price data are obtained from Thomson Reuters Datastream for both spot and front month futures prices.[13] All observations of feed ingredient prices are converted to NOK with an underlying assumption that salmon farmers already have a perfect exchange rate hedge in place. We apply fixed exchange rates of 7.0140 NOK/USD and 8.6732 NOK/EUR. 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 Fig. 1. Salmon is depicted in the top panel and soymeal, wheat and rapeseed oil in the

[13] Exact name and ticker for each price series obtained are listed in Appendix B.1. Price series that are denoted in bushels or tons are converted to kilograms.

Table 1
Descriptive statistics for in-sample weekly spot and futures returns of salmon, soymeal, wheat and rapeseed oil.

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 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.

lower panel.

Each time series consist of 643 observations collected from January 2008 to April 2020. For the purpose of our modeling, we use the log-transformed percentage return[14] series which are presented in Appendix B.1. 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. Descriptive statistics for the in-sample period are presented in Table 1.[15] Results show that all bar one of the returns distributions differ significantly from the normal distribution. 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. This distinction between positive and negative returns will be important in discussions of results in Section 5.

Several tests were conducted to verify the suitability of GARCH models, which are presented in Table 2.[16] Test results strongly reject the null hypotheses of a unit root for any of the returns series and are stable across different lag lengths.[17] A visual inspection of the ACF and PACF plots found in Appendix B suggest that the AR models capture the autocorrelation well. The Ljung-Box Q test (LBQ) show no evidence of autocorrelation.[18] The Engle's Lagrange multiplier (LM) test strongly reject the null hypotheses of no ARCH effects.[19] 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.

4.2. Estimated models

In this section, we present estimation results for the uni- and multivariate models described in Section 2 and illustrate the resulting one-ahead rolling window forecasts for conditional standard deviation.

We find that the estimates for the univariate GARCH(1,1) models 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. This can be seen by visual examination of the salmon futures log-return plot in Appendix B.1. An exception is 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 $\hat{\alpha} = 0$, hence we elect to keep the model. Furthermore, we test the models for miss-specification by the weighted ARCH LM test (Fisher and Gallagher, 2012), the

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

[15] Descriptive statistics for the out-of-sample period can be found in Appendix B.1.

[16] All quoted test statistics are for tests with no drift and no trend. 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 $k = \text{int}\left(12(T/100)^{1/4}\right)$, where T denotes sample size.

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

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

[19] Tests performed with 12 lags.

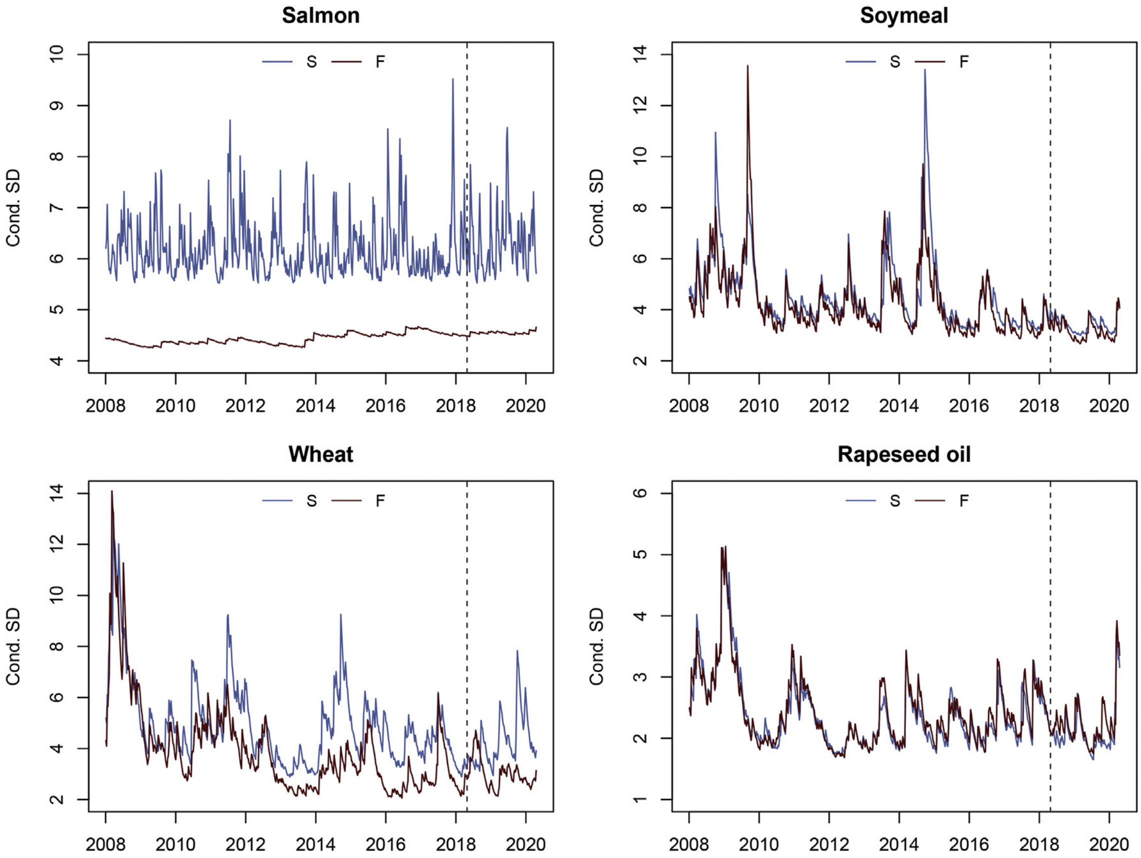

Fig. 2. One-ahead rolling forecasts for conditional standard deviation, σ_t.

Nyblom stability test (Nyblom, 1989) and the adjusted Pearsons goodness-of-fit test (Vlaar and Palm, 1993). Results indicate that the models are adequately specified.[20]

Fig. 2 presents the one-ahead rolling forecast of conditional standard deviation for the estimated GARCH(1,1) models. The dashed vertical line separates the in- and out-of-sample period. We observe 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 feature substantial differences, where the salmon spot is volatile and changing frequently, with low long term volatility persistence (i.e. volatility spikes fade quickly as $(\widehat{\alpha} + \widehat{\beta}) \ll 1$).

Lastly, we estimate parameters for the multivariate models (DCC, C-DCC, TVC and RWC). Estimates and further comments are attached in Appendix B.3. When estimating the TVC model, we elect to solely model the spot and futures pair for each respective commodity, implying that we solely estimate the single-hedge. 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. The lag length is set to $n = 1$ as this during estimation has shown to yield the most reasonable models. Wu (2018) also uses lag length $n = 1$ when hedging grain sorghum and finds reasonable results.

When estimating the RWC model, there is limited guidance on how to select the most appropriate rolling window size. 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. Although 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.

5. Results

We begin by presenting dynamic hedge ratios obtained by the models proposed in Section 2. Table 3 shows in-sample mean, minimum, maximum and standard deviation (SD) for each model and commodity.[21] 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

[20] Further details on tests and estimates are attached in Appendix B.2.

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

Table 3
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: Results for the dynamic conditional correlation model (DCC), copula dynamic conditional correlation model (C-DCC), time-varying copula model (TVC) and the rolling window copula model (RWC). Smallest and largest standard deviation (SD) for each commodity are marked red and blue respectively.

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 Figs. 3 and 4, 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 portfolios require fewer adjustments between each time-step than the latter.[22] There is considerably higher SD for the TVC model in soymeal, which is a result of extreme negative spikes in Fig. 4. Potentially, a result of extreme tail events which are more accurately captured by the copula models.[23] For the same reason, TVC is the only model with a substantial 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).

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 favorable 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 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

[22] 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.

[23] The soymeal spot and futures pair is modeled by the Student's t TVC model, which emphasizes tail dependence.

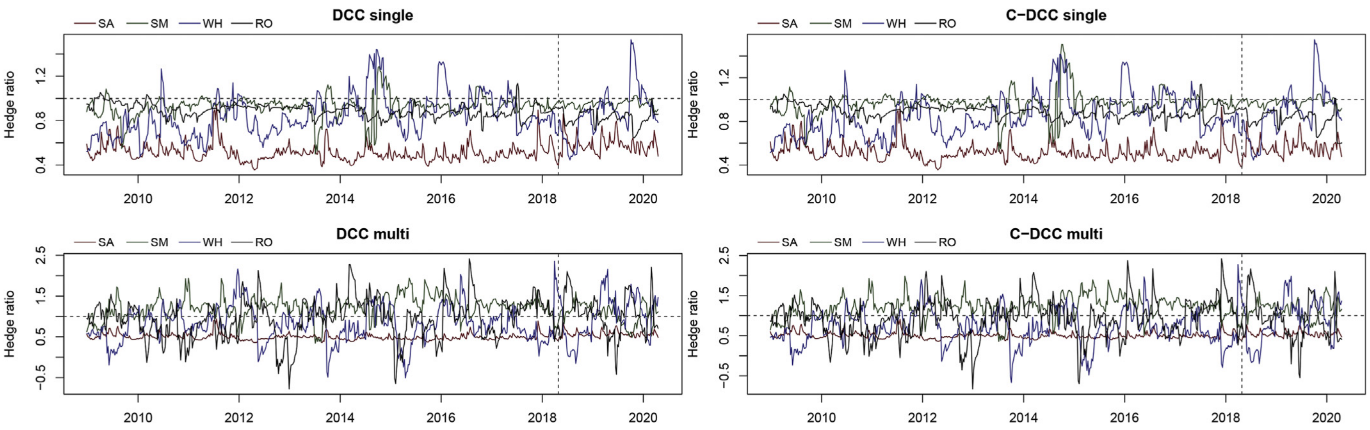

(a) DCC **(b)** C-DCC

Fig. 3. Hedge ratios for the dynamic conditional correlation (DCC) models and the copula-DCC (C-DCC) models.

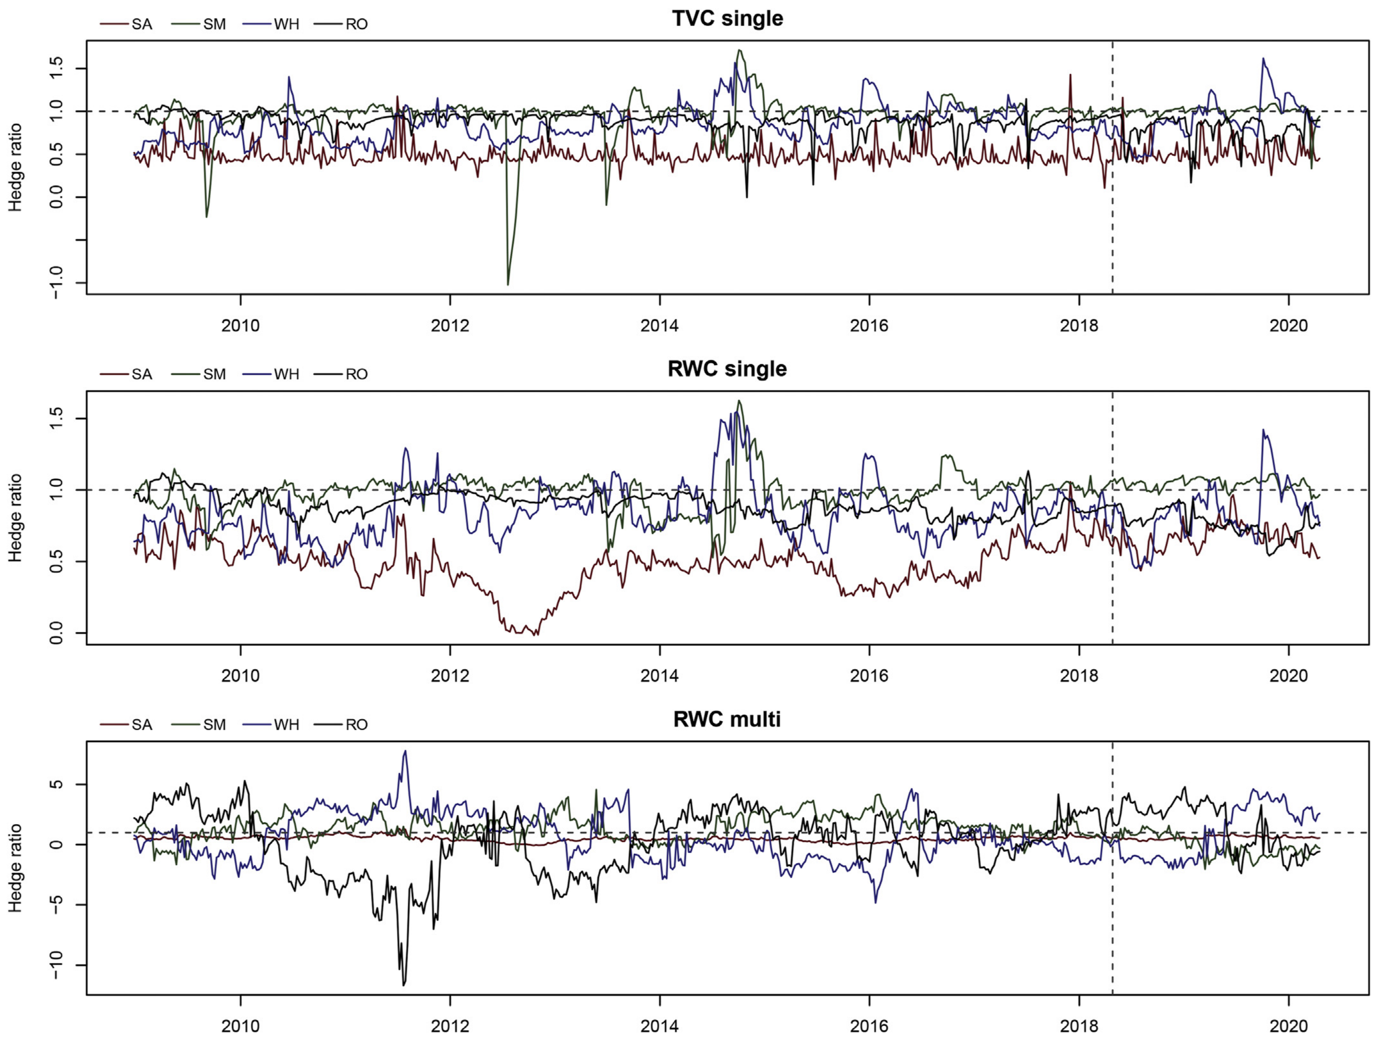

Fig. 4. Hedge ratios for the time-varying copula (TVC) model and rolling window copula (RWC) models.

Table 4
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, for the dynamic conditional correlation model (DCC), copula dynamic conditional correlation model (C-DCC), time-varying copula model (TVC) and the rolling window copula model (RWC). Smallest and largest standard deviation (SD) for each commodity are marked red and blue respectively. Lowest and highest values for mean return and hedge effectiveness are marked as red and blue respectively.

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*,[24] 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).[25] 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 summarize, the modeling extension of applying copulas tends to increase the standard 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.

5.1. Hedging outcomes

Hedging outcomes are the result of applying Equation (14) to the price data of Section 4.1 and hedge ratios of Table 3 and C.13. The results are a distribution of hedging outcomes, which are summarized by the performance measures of Section 2. Table 4 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.

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 (1046 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).[26] 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 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 3 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

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

[25] 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.

[26] 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.

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.[27] 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%)[28] is greater than for the single-hedge (23.70%, 26.13%).[28] 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, 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 earlier. 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. Notice that the unhedged portfolio for the out-of-sample period has a *negative* mean return of −2943 NOK. This implies a period of overall price depreciation as described in Section 4.1, 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 optimize by using the RWC model during good periods and the naïve strategy during bad. Furthermore, we observe that the RWC model, which performed 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, −1306 NOK mean return (−1 888, −1972 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 data sensitive due to the specification of (13), also noted by Patton (2006a). Consequently, the model does not generalize well out-of-sample.

Fig. 5 shows the distributions of returns for different hedging strategies for the in-sample period.[29] 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 piece-wise 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 Fig. 5b. 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 center). 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). Fig. 5 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 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.

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 favor dynamic hedging models. This is confirmed in Table 5, 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

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

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

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

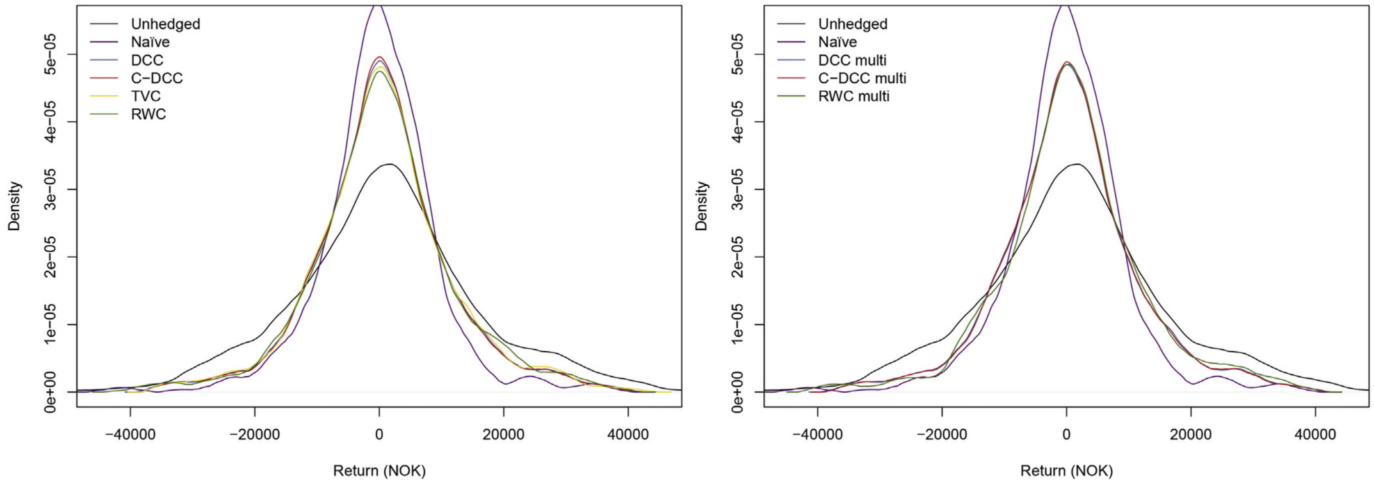

(a) Static and dynamic single-hedges.

(b) Static and dynamic multi-hedges.

Fig. 5. In-sample return distributions for the four-week hedging horizon. Models depicted are the dynamic conditional correlation model (DCC), copula dynamic conditional correlation model (C-DCC), time-varying copula model (TVC) and the rolling window copula model (RWC).

Table 5
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, for the dynamic conditional correlation model (DCC), copula dynamic conditional correlation model (C-DCC), time-varying copula model (TVC) and the rolling window copula model (RWC). Lowest and highest values for mean return and hedge effectiveness are marked as red and blue respectively.

characteristics and sample period.[30]

To summarize, 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 (−1177) 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 extreme 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 favorable in terms of HE and costs, but requires pre-planned slaughtering volumes to a higher degree.

5.2. Sensitivity to hedging horizon

To analyze the sensitivity to different horizons, we provide results for horizons between 1 and 20 weeks for the full sample.[31] Fig. 6a 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 $n = 1$.[32] 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 5.1. Fig. 6b clearly illustrates this, showing mean return adjusted for transaction costs for different horizons. 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 hedge performs the worst, due to the low MR and high transaction costs.

5.3. Cost of hedging

In this section, we derive the implications for the trade-off between risk and return. We do this by proposing a measure of the cost 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- and hedged portfolio. 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

30 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.

31 In this subsection, we elect to use the full sample to focus on the sensitivity and not the specific differences between in- and out-of-sample.

32 Modeling wise, there is limited guidance for selecting an appropriate lag length (Patton, 2006a).

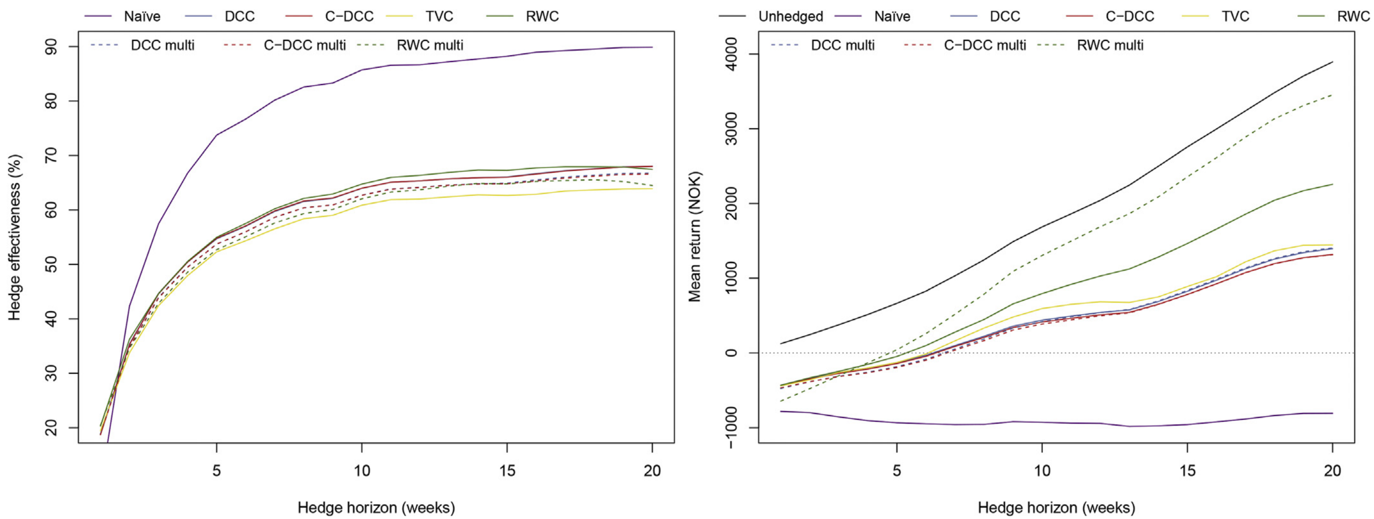

(a) Hedge effectiveness. **(b)** Transaction cost adjusted mean return.

Fig. 6. Hedging performance for different hedging horizons, full sample. Models depicted are the dynamic conditional correlation model (DCC), copula dynamic conditional correlation model (C-DCC), time-varying copula model (TVC) and the rolling window copula model (RWC).

Table 6
Hedging cost measures for the 4-week hedging horizon.

	Unhedged	Naïve	DCC		C-DCC		TVC	RWC	
			Single	Multi	Single	Multi	Single	Single	Multi
In-sample (*N*= 538)									
Mean return, TCA	1046	−777	62	13	54	14	99	159	186
CoH	0	1823	984	1033	992	1032	947	887	860
HE (%)	NA	64.39%	49.04%	48.15%	49.84%	48.91%	47.70%	46.00%	45.10%
CVR	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	−2943	−1413	−1945	−2020	−2016	−2101	−2110	−1888	−1972
CoH	0	−1530	−997	−923	−927	−842	−832	−1055	−971
HE (%)	NA	70.07%	52.24%	51.25%	50.95%	49.93%	47.66%	57.24%	53.52%
CVR	NA	−21.83	−19.09	−18.01	−18.19	−16.86	−17.47	−18.44	−18.15

Note: Results denoted in NOK, for the dynamic conditional correlation model (DCC), copula dynamic conditional correlation model (C-DCC), time-varying copula model (TVC) and the rolling window copula model (RWC). TCA denotes transaction cost adjusted returns.

of a given hedging strategy. A natural extension is the ratio $\frac{CoH}{HE}$, which is the cost per percentage of variance reduction, or cost of variance reduction (CVR) for short. A lower CVR is preferred, in contrast to hedge effectiveness. 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.

Second, we compare the CVR of models for the four-week horizon as presented in Table 6. Recall from the previous discussion on the risk-return trade-off that our results indicate a non-linear relationship between hedge effectiveness and mean return. If the relationship was linear, we would expect the CVR to be the same for all models. An increase in HE would then 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 CVR for the respective models, as shown by Table 6. For the in-sample period, we find the RWC multi-hedge yields the best CVR (19.07), while the naïve hedge yields the worst (28.31), a point difference of 9.24. All the dynamic models perform similarly, with CVR falling in the range from 19.07 (RWC multi) to 21.45 (DCC multi). In general, we see that cost of variance reduction 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 CVR (−21.83) and the C-DCC multi-hedge give the worst CVR (−16.86). This also conforms 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. In other words, low reduction of mean return gains during periods of positive return should imply low reduction of mean return *losses* during periods of negative return. On 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 CVR to hedging horizon. Fig. 7 shows CVR 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 the multi-hedge. This can be seen in Fig. 6b, 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 CVR 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 Fig. 7b 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 summarize, these results confirm our previous discussions. Although the naïve hedge is superior in terms of hedging effectiveness and for periods of negative returns, it tends to over-hedge. This means that 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.

6. Conclusion

In this paper, we address the price risk hedging problem for farmers of Atlantic salmon. We analyze 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.

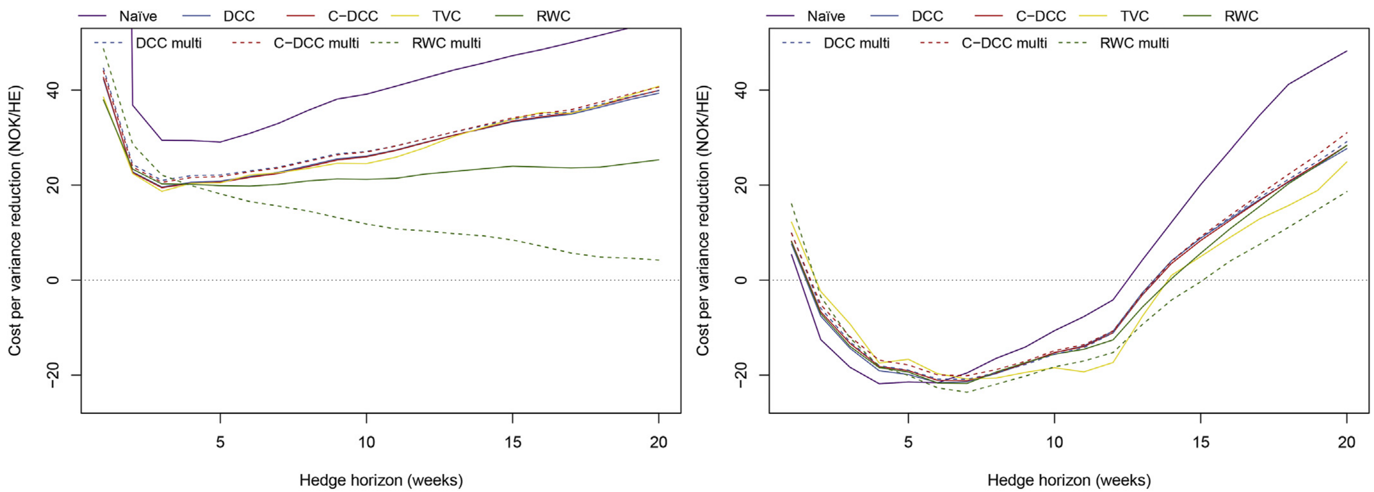

(a) In-sample. **(b)** Out-of-sample.

Fig. 7. Cost of variance reduction for different hedging horizons. Models depicted are the dynamic conditional correlation model (DCC), copula dynamic conditional correlation model (C-DCC), time-varying copula model (TVC) and the rolling window copula model (RWC).

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 summarized 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 that sacrifices the least return per unit 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 favorable in terms of hedge effectiveness and costs, but requires pre-planning of slaughtering volumes to a higher degree.

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

In what follows we propose several promising directions for further research. First, a promising extension of the current framework would be to explicitly account for liquidity issues on salmon contracts, and find optimal entering and exit strategies. Second, future research could further investigate how different samples and sample sizes affect the hedging performance. Lastly, it would be interesting to further explore the connection between the input commodity markets and the hedging results, by analyzing how the inclusion of input commodities in the multi-hedge contributes to improved results over longer hedging horizons.

Author statement

Aleksander H. Haarstad: Conceptualization, Methodology, Writing- Original draft preparation, Writing - Reviewing and Editing. Maria Lavrutich: Conceptualization, Supervision, Writing- Reviewing and Editing. Kristian Strypet: Formal analysis, Writing- Original draft preparation. Eivind Strøm: Methodology, Software, Investigation, Writing- Original draft preparation.

Declaration of competing interest

There are no financial conflicts of interest to disclose.

Acknowledgments

We thank the Editor, Betty Simkins, as well as two anonymous referees and associate editor, for their thoughtful comments. We also thank Kåre Gruven, Chief Feed Adviser at Norway Royal Salmon, for his valuable insights.

Appendix. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.jcomm.2021.100182.

References

Aas, T.S., Ytrestøyl, T., Åsgård, T., 2019. Utilization of feed resources in the production of Atlantic salmon (Salmo salar) in Norway: an update for 2016. Aquaculture Reports 15, 100216, https://doi.org/10.1016/j.aqrep.2019.100216.

Abolofia, J., Asche, F., Wilen, J.E., 2017. The cost of lice: quantifying the impacts of parasitic sea lice on farmed salmon. Mar. Resour. Econ. 32, https://doi.org/10.1086/691981.

Acerbi, C., Tasche, D., 2001. On the coherence of expected shortfall. J. Bank. Finance 26, 1487–1503, https://doi.org/10.1016/S0378-4266(02)00283-2.

Aepli, M.D., Füss, R., Henriksen, T.E.S., Paraschiv, F., 2017. Modeling the multivariate dynamic dependence structure of commodity futures portfolios. Journal of Commodity Markets 6, 66–87, https://doi.org/10.1016/j.jcomm.2017.05.002.

Andersen, B.P., 2019. Efficiency in the atlantic salmon futures market. https://hdl.handle.net/11250/2622898.

Anderson, R.W., Danthine, J.P., 1981. Cross hedging. J. Polit. Econ. 89, 1182–1196, https://doi.org/10.1086/261028.

Anderson, D.P., Fei, C., Stevens, R., Vedenov, D., 2017. Multi-commodity hedging in the live cattle futures market. Agricultural & Applied Economics Association, https://doi.org/10.22004/ag.econ.258560.

Asche, F., 2017. New markets, new technologies and new opportunities in aquaculture. Aquacult. Econ. Manag. 21, https://doi.org/10.1080/13657305.2016.1272649.

Asche, F., Guttormsen, A.G., Tveteras, R., 2008. Aquaculture—opportunities and challenges. Mar. Resour. Econ. 23, 395–400.

Asche, F., Misund, B., Oglend, A., 2016a. Determinants of the Atlantic salmon futures risk premium. Journal of Commodity Markets 2, 6–17, https://doi.org/10.1016/j.jcomm.2016.07.001.

Asche, F., Misund, B., Oglend, A., 2016b. The spot-forward relationship in the Atlantic salmon market. Aquacult. Econ. Manag. 20, 222–234, https://doi.org/10.1080/13657305.2016.1156192.

Basu, P., Gavin, W.T., 2017. Negative correlation between stock and futures returns: an unexploited hedging opportunity? Bull. Econ. Res. 69, 209–215, https://doi.org/10.1111/boer.12090.

BioMar Group, 2019. Annual Report 2018. https://www.biomar.fr/globalassets/.global/pdf-files-_en/biomar-group-a nnual-report-2018.pdf. (Accessed 11 October 2019).

Bloznelis, D., 2018. Hedging salmon price risk. Aquacult. Econ. Manag. 22, 168–191, https://doi.org/10.1080/13657305.2017.1409292.

Brakstad, O.M., Hagspiel, V., Lavrutich, M.N., Matanovic, D., 2019. Optimal investment decisions in lice-fighting technologies: a case study in Norway. Aquaculture 504, 300–313, https://doi.org/10.1016/j.aquaculture.2019.01.040.

Brooks, C., 2014. Introductory Econometrics for Finance, third ed. Cambridge University Press (CUP.

Butterworth, D., Holmes, P., 2001. The hedging effectiveness of stock index futures: evidence for the FTSE-100 and FTSE-mid250 indexes traded in the UK. Appl. Financ. Econ. 11, 57–68, https://doi.org/10.1080/09603100150210264.

CME Group, 2014. Agricultural commodity metric conversion guide. https://www.cmegroup.com/trading/agricultural/files/AC-225_MetricGuide. pdf.

CME Group, 2020. Exchange fees for clearing & trading. https://www.cmegroup.com/company/clearing-fees.html, (Accessed 29 May 2020).

Dahl, R.E., Jonsson, E., 2018a. Volatility spillover in aquaculture and fisheries markets. Aquacult. Econ. Manag. 22, https://doi.org/10.1080/13657305.2018.1429033.

Dahl, R.E., Jonsson, E., 2018b. Volatility spillover in seafood markets. Journal of Commodity Markets 12, https://doi.org/10.1016/j.jcomm.2017.12.005.

Dahl, R.E., Oglend, A., 2014. Fish price volatility. Mar. Resour. Econ. 29, 305–322, https://doi.org/10.1086/678925.

Dahl, R.E., Yahya, M., 2019. Price volatility dynamics in aquaculture fish markets. Aquacult. Econ. Manag. 23, 321–340, https://doi.org/10.1080/13657305.2019.1632390.

Ederington, L.H., 1979. The hedging performance of the new futures markets. J. Finance 34, 157, https://doi.org/10.2307/2327150.

Engle, R., 2002. Dynamic conditional correlation: a simple class of multivariate generalized autoregressive conditional heteroskedasticity models. J. Bus. Econ. Stat. 20, 339–350, https://doi.org/10.1198/073500102288618487.

Fackler, P.L., McNew, K.P., 1993. Multiproduct hedging: theory, estimation, and an application. Rev. Agric. Econ. 15, 521, https://doi.org/10.2307/1349486.

Fama, E.F., 1990. Stock returns, expected returns, and real activity. J. Finance 45, 1089, https://doi.org/10.2307/2328716.

Fama, E.F., French, K.R., 1992. The cross-section of expected stock returns. J. Finance 47, 427–465, https://doi.org/10.1111/j.1540-6261.1992.tb04398.x.

Fernández, C., Steel, M.F., 1998. On bayesian modeling of fat tails and skewness. J. Am. Stat. Assoc. 93, 359–371, https://doi.org/10.1080/01621459.1998.10474117.

Fish Pool, 2016. Fish Pool index. https://fishpool.eu/price-information/spot-prices/, (Accessed 11 November 2019).

Fish Pool, 2020. Fee list. http://fishpool.eu/trading/fee-list/, (Accessed 25 May 2020).

Fisher, T.J., Gallagher, C.M., 2012. New weighted portmanteau statistics for time series goodness of fit testing. J. Am. Stat. Assoc. 107, 777–787, https://doi.org/10.1080/01621459.2012.688465.

Franken, J.R., Parcell, J.L., 2011. Cross-hedging fishmeal: exploring corn and soybean meal futures contracts. Aquacult. Econ. Manag. 15, 71–81, https://doi.org/10.1080/13657305.2011.549166.

French, K.R., Fama, E.F., 1989. Business conditions and expected returns on stocks and bonds. J. Financ. Econ. 25, 23–49, https://doi.org/10.1016/0304-405X(89)90095-0.

Ghalanos, A., 2019. The Rmgarch Models : Background and Properties. https://cran.r-project.org/web/packages/rmgarch/vignettes/The_rmgarch_m odels.pdf. Version 1.3-0(Accessed 24 October 2019).

Graham, J.R., Smith, C.W., 1999. Tax incentives to hedge. J. Finance 54, 2241–2262, https://doi.org/10.1111/0022-1082.00187.

Hagspiel, V., Hannevik, J., Lavrutich, M., Naustdal, M., Struksnæs, H., 2018. Real options under technological uncertainty: a case study of investment in a post-smolt facility in Norway. Mar. Pol. 88, 158–166, https://doi.org/10.1016/j.marpol.2017.11.020.

Haigh, M.S., Holt, M.T., 2000. Hedging multiple price uncertainty in international grain trade. Am. J. Agric. Econ. 82, 881–896, https://doi.org/10.1111/0002-9092.00088.

Hansen, P.R., Lunde, A., 2005. A forecast comparison of volatility models: does anything beat a GARCH(1,1)? J. Appl. Econom. 20, 873–889. https://www.jstor.org/stable/25146403.

Iversen, A., Asche, F., Hermansen, Ø., Nystøyl, R., 2020. Production cost and competitiveness in major salmon farming countries 2003–2018. Aquaculture 522, 735089, https://doi.org/10.1016/j.aquaculture.2020.735089.

Jondeau, E., Rockinger, M., 2006. The Copula-GARCH model of conditional dependencies: an international stock market application. J. Int. Money Finance 25, 827–853, https://doi.org/10.1016/j.jimonfin.2006.04.007.

Lerøy Seafood Group, 2020. Annual report 2019. https://www.leroyseafood.com/en/investor/reports-and-webcast/annual-rep ort-2019/, (Accessed 25 May 2020).

Misund, B., Asche, F., 2016. Hedging efficiency of Atlantic salmon futures. Aquacult. Econ. Manag. 20, 368–381, https://doi.org/10.1080/13657305.2016.1212123.

Mowi, 2019. Salmon farming industry handbook 2019. https://corpsite.azureedge.net/corpsite/wp-content/uploads/2019/06/Salm on-Industry-Handbook-2019.pdf, (Accessed 2 October 2019).

Mowi, 2020. Integrated annual report 2019. https://corpsite.azureedge.net/corpsite/wp-content/uploads/2020/03/Mowi _Annual_Report_2019.pdf, (Accessed 1 May 2020).

Nagler, T., Czado, C., 2016. Evading the curse of dimensionality in nonparametric density estimation with simplified vine copulas. J. Multivariate Anal. 151, 69–89, https://doi.org/10.1016/j.jmva.2016.07.003 arXiv:1503.03305.

Norges Bank, 2020. Valutakurser. https://www.norges-bank.no/tema/Statistikk/Valutakurser/, (Accessed 24 May 2020).

Nyblom, J., 1989. Testing for the constancy of parameters over time. J. Am. Stat. Assoc. 84, 223–230, https://doi.org/10.1080/01621459.1989.10478759.

Oglend, A., 2013. Recent trends in salmon price volatility. Aquacult. Econ. Manag. 17, 281–299, https://doi.org/10.1080/13657305.2013.812155.

Oglend, A., Straume, H.M., 2020. Futures market hedging efficiency in a new futures exchange: effects of trade partner diversification. J. Futures Mark. 40, 617–631, https://doi.org/10.1002/fut.22088.

Veikkaus Oy, 2015. Sustain ability. In: The Strategy Pathfinder. John Wiley & Sons, Inc., Hoboken, NJ, USA, pp. 301–335, https://doi.org/10.1002/9781119206026.ch10.

Parcell, J., Pierce, V., 2011. Agricultural Commodity Futures Contract Specifications. Department of Agricultural and Applied Economics, Columbia. Technical Report https://extension2.missouri.edu/G601.

Patton, A.J., 2006a. Estimation of multivariate models for time series of possibly different lengths. J. Appl. Econom. 21, 147–173, https://doi.org/10.1002/jae.865.

Patton, A.J., 2006b. Modelling asymmetric exchange rate dependence∗. Int. Econ. Rev. 47, 527–556, https://doi.org/10.1111/j.1468-2354.2006.00387.x.

Patton, A.J., 2012. A review of copula models for economic time series. J. Multivariate Anal. 110, 4–18, https://doi.org/10.1016/j.jmva.2012.02.021.

Pincinato, R.B., Asche, F., Bleie, H., Skrudland, A., Stormoen, M., 2021. Factors influencing production loss in salmonid farming. Aquaculture 532, https://doi.org/10.1016/j.aquaculture.2020.736034.

Power, G.J., 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. https://legacy.farmdoc.illinois.edu/nccc134/conf_2008/pdf/confp11-08.pdf.

Power, G.J., Vedenov, D., 2009. Dealing with downside risk in a multi-commodity setting: a case for a “Texas hedge”? J. Futures Mark. 30, 290–304, https://doi.org/10.1002/fut.20411.

Power, G.J., Vedenov, D., Anderson, D., Klose, S., 2013. Market volatility and the dynamic hedging of multi-commodity price risk. Appl. Econ. 45, 3891–3903, https://doi.org/10.1080/00036846.2012.736942.

Rosenberg, J.V., Schuermann, T., 2006. A general approach to integrated risk management with skewed, fat-tailed risks. J. Financ. Econ. 79, 569–614, https://doi.org/10.1016/j.jfineco.2005.03.001.

SalMar, 2020. Annual report 2019. https://ml-eu.globenewswire.com/Resource/Download/74e5d277-5302-41ee-af 23-4a7785858807, (Accessed 4 May 2020).

Schmidbauer, H., Roesch, A., Tunalioglu, V.S., 2016. mgarchBEKK: simulating, estimating and diagnosing MGARCH (BEKK and mGJR) processes. https://cran.r-project.org/package=mgarchBEKK. r package version 0.0.2.

Schütz, P., Westgaard, S., 2018. Optimal hedging strategies for salmon producers. Journal of Commodity Markets 12, 60–70, https://doi.org/10.1016/j.jcomm.2017.12.009.

Schwert, G.W., 2002. Tests for unit roots: a Monte Carlo investigation. J. Bus. Econ. Stat. 20, 5–17, https://doi.org/10.1198/073500102753410354.

Sklar, A., 1959. Fonctions de répartition à n dimensions et leurs marges. Publications de l’Institut Statistique de l’Université de Paris, pp. 229–231.

Smith, C.W., Stulz, R.M., 1985. The determinants of firms’ hedging policies. J. Financ. Quant. Anal. 20, 391–405, https://doi.org/10.2307/2330757.

Steen, M., Jacobsen, F., 2020. Modeling the return distribution of salmon farming companies: a quantile regression approach. Aquacult. Econ. Manag. 24, 310–337, https://doi.org/10.1080/13657305.2020.1765896.

Thyholdt, S.B., 2014. The importance of temperature in farmed salmon growth: regional growth functions for Norwegian farmed salmon. Aquacult. Econ. Manag. 18, 189–204, https://doi.org/10.1080/13657305.2014.903310.

Toft, K.B., 1996. On the mean-variance tradeoff in option replication with transactions costs. J. Financ. Quant. Anal. 31, 233–263, https://doi.org/10.2307/2331181.

Tse, Y.K., Tsui, A.K.C., 2002. A multivariate generalized autoregressive conditional heteroscedasticity model with time-varying correlations. J. Bus. Econ. Stat. 20, 351–362, https://doi.org/10.1198/073500102288618496.

Vlaar, P.J., Palm, F.C., 1993. The message in weekly exchange rates in the European monetary system: mean reversion, conditional heteroscedasticity, and jumps. J. Bus. Econ. Stat. 11, 351–360, https://doi.org/10.1080/07350015.1993.10509963.

Vukina, T., Anderson, J.L., 1993. A state-space forecasting approach to optimal intertemporal cross-hedging. Am. J. Agric. Econ. 75, 416–424, https://doi.org/10.2307/1242926.

Wu, S., 2018. Cross-hedging with agricultural commodities: a copula-based GARCH model. http://shenanwu.wordpress.ncsu.edu/files/2018/05/Cross-hedging-with-Agricultural-Commodities2.pdf, (Accessed 9 November 2019).

Zhao, J., Goodwin, B.K., 2012. Dynamic cross-hedge ratios: an application of copula models. https://ageconsearch.umn.edu/record/124610/files/AAEA 2012 Dynamic Cross-Hedge Ratios_ Zhao_ 1.pdf, (Accessed 3 November 2019).