ISBN: 978-82-575-1377-1
ISSN: 1894-6402

Norwegian University
of Life Sciences

Postboks 5003
NO-1432 Ås, Norway
+47 67 23 00 00
www.nmbu.no

Philosophiae Doctor (PhD), Thesis 2016:56

Daumantas Bloznelis

Norwegian University of Life Sciences
Faculty of Social Sciences
School of Economics and Business

Philosophiae Doctor (PhD)
Thesis 2016:56

Management of short-term price uncertainty in the salmon spot market

Styring av kortsiktig prisusikkerhet i laksespotmarkedet

Daumantas Bloznelis

Management of short-term price uncertainty in the salmon spot market

Styring av kortsiktig prisusikkerhet
i laksespotmarkedet

Philosophiae Doctor (PhD) Thesis

Daumantas Bloznelis

School of Economics and Business
Faculty of Social Sciences
Norwegian University of Life Sciences

Ås/Adamstuen (2016)

Thesis number 2016:56
ISSN 1894-6402
ISBN 978-82-575-1377-1

Acknowledgements

A number of people have kindly helped me in the process of writing the thesis, and I am very grateful for their contributions. I thank my supervisor prof. Atle Guttormsen for granting me ultimate intellectual freedom, providing financial independence via the Ph.D. contract and enabling me to visit Cornell University in the U.S., where I tasted and enjoyed the academic life of the Ivy League. I thank prof. Ole Gjølberg for voluntarily undertaking the hard and high-responsibility job as a supervisor; for directing me gently to interesting research topics; for his endless optimism, energy and encouragement; and for his scrupulous and rigorous editing of my first draft ("send me an electronic version – there is not enough space for my comments on the paper"). I thank all the participants at the *scientific coffee* and *scientific lunch* meetings; life would always look brighter on Thursday afternoons after the *scientific coffee*, and work would go slow on Mondays if not for the *scientific lunch*. I thank prof. Marie Steen for her continuous encouragement and genuine interest in my progress; for her question "What is your contribution?"; and for constructive comments on my drafts. Prof. Olvar Bergland and prof. Kyrre Rickertsen were there for me to discuss econometrics, which I highly appreciate. I thank prof. William Warner for teaching me the principles of academic writing and showing me how to plan and write in a structured and efficient way. His method has become the backbone of my writing and research planning, and I do not hesitate to advertise and recommend it to my peers. A similar thanks is due for prof. Stephen Mumford for the "Mumford method"; for "think slow; write fast"; and for the concept of the "Mumford sheet". Berit Pettersen helped me cope with challenges regarding funding, taxes and bureaucracy with outstanding style and admirable ease. Reidun Aasheim could always be trusted when facing any practical problems at work. Both of them saved me from a lot of unnecessary stress and frustration.

My gratitude extends beyond the professors and staff at NMBU. I thank prof. Harry Kaiser for inviting me to Cornell University and being a great teacher and a welcoming host. I thank prof. Yongmiao Hong for his lectures of unmatched clarity; for his patience answering my questions on asymptotic statistics and econometrics; and for great conversations and delicious food at his Sunday seminars. I thank prof. George Jakubson for explaining the principles of econometrics in the most entertaining way I could have ever dreamt of; for his kindness answering my questions that would take him hours long; and for his brilliant maxims on econometric practice. Prof. Jūratė Liaukonytė was of immense help at various stages of my progress, which I am grateful for. I thank prof. Timo Teräsvirta for his time spent sharing insights regarding volatility modelling. I thank prof. Ruey S. Tsay for his great "Advanced Time Series Analysis" course that provided me a new angle for viewing time series econometrics at; and for his authoritative answers to questions not encountered in textbooks.

I thank prof. Remigijus Leipus and prof. Virmantas Kvedaras for their dedication helping me with econometric modelling. I thank prof. Alfredas Račkauskas for supporting my academic career the whole way from the bachelor studies to the Ph.D. by challenging and encouraging, and always being optimistic and very friendly.

I thank prof. Per Bjarte Solibakke, prof. Gabriel Power, prof. Gunnar Bårdsen and prof. Ronald Huisman for being great discussants at the conferences and providing insightful, constructive and encouraging comments on my manuscripts.

I thank Inger Egebakken and Berit Søvik for teaching me the Norwegian language and helping me prepare for and pass the "Test of Norwegian – advanced level", also known as "Bergenstesten". Being able to communicate in Norwegian helped me open new doors and attract new colleagues to collaborate with.

Life would have been miserable without my fellow Ph.D. students. I thank Erik Smith-Meier, my office colleague, for becoming my first and best friend in Norway; for advising me on various aspects of life in Norway and helping me settle down; and for his tolerance of my veering between econometric considerations and personal feelings, both positive and negative. I thank Lin Ma, Arnar Mar Buason and Andrej Stenšin, my fellow Ph.D. students and office colleagues, for tolerating me in the office and being friendly and kind. I thank Senyonga Livingstone, Xianwen Chen and Arnar Mar Buason (again) for interesting discussions about econometrics without which I would have felt intellectually isolated. The couple of past-midnight hours spent walking in the cold around a garage block with Xianwen, discussing whether it is possible to infer predictive power from in-sample modelling, were indeed a precious and genuine act of science in the making. I thank Torun Sæther Fretheim for showing compassion during a difficult time and bringing me into the commodities research group; it was a key turning point that got this thesis going. I thank Ruitong Wang for econometric conversations at Cornell University and for fetching research papers from databases I did not have access to.

I thank the members of the Cross Validated and Stack Overflow online communities for their altruistic attitude and patience in helping me, as well literally millions of other people, understand statistics and solve particular statistical and programming problems. I thank Ås karate klubb (especially sensei Birger Kristoffersen, sempai Cathrine Glosli and sempai Anja Velure) for keeping me fit and energized, and for providing a friendly environment.

Last but not the least, I thank my aunt Laima Bloznelytė-Plėšnienė for *scaring* me into taking a MSc degree, which was a necessary prerequisite for the Ph.D. studies. I thank my father prof. Mindaugas Bloznelis for educating me in mathematics all the way from kindergarten to the Ph.D.; for providing deep insights into foreign territories of statistics, econometrics and beyond – without him most of my econometric techniques and understanding would be floating bellies up in the river of ignorance; and for being a personal example of no-compromise hard work and perseverance. I thank my mother Eglė Blozneliene for supporting me firmly and accepting my ups and downs without judging; she was there for me when I needed her the most. I thank my grandmother Eugenija Stravinskienė for having taught me not to fear large-scale projects but rather work methodically, bit by bit, until even the biggest job is done. I thank my sister Kotryna Bloznelytė for digging up research papers I could not get in the local library, and for consulting me regarding career opportunities. I cannot thank enough my dear wife Dovilė Blozneliene for her unconditional love, care and support, and endless patience that I did not know was humanly possible. She was with me every day for these challenging years and she *never ever* turned her back on me.

More acknowledgements are certainly due, but not all can be delivered due to space limitations and my porous memory. Those not mentioned here should not feel insulted; if you have helped me, I am indeed very grateful!

Ås, February 2016

Daumantas Bloznelis

Abstract

The spot price of Norwegian farmed Atlantic salmon is highly volatile and hard to predict. The uncertainty over the future spot price obscures revenue projections of salmon farmers and cost expectations of processors, exporters and retailers. It also makes business financing expensive as the high uncertainty needs to be compensated by a high return on investment. The market participants acknowledge this to be a substantial quandary. This thesis examines the problem and aims to provide feasible solutions for uncertainty management in the salmon market.

An introduction and three research papers address the different aspects of the subject, namely, salmon price volatility, price predictability and hedging the spot price with various financial instruments. A variety of econometric and machine learning techniques are applied to account for seasonal patterns and autoregressive conditional heteroskedasticity in the price series and to deliver forecasts of their conditional means and variances. The first paper "Salmon price volatility: a weight-class-specific multivariate approach" presents a statistical description of the conditional mean and variance of the spot prices of seven different weight classes of salmon. It highlights a considerable increase in the unconditional variance around 2006, which coincides with a change in industry regulations and the introduction of a futures exchange for salmon. The conditional mean and variance patterns are found to be similar across the neighbouring weight classes, and the conditional correlations are nearly perfect since 2007. This allows treating the three most popular weight classes of 3-4 kg, 4-5 kg and 5-6 kg as one and makes hedging with salmon futures relatively attractive. The second paper "Short-term salmon price forecasting" is a comprehensive study of forecasting the spot price one to five weeks ahead. It employs three different classes of forecasting models: (1) time series models broadly based on the ARIMA model; (2) artificial neural networks; and (3) a custom model based on the k-nearest neighbours method. Six measures of forecast accuracy and seven tests of forecast optimality and encompassing are reported. The salmon price appears to have a partly predictable seasonal component; however, statistical significance of predictability cannot be established at the available sample size, and the economic value of forecasts is limited. Also, unpredictability beyond seasonality does not offer evidence against weak form efficiency of the salmon spot market. The third paper "Hedging salmon price risk" examines the hedging performance of salmon futures, live cattle futures, soybean meal and oil futures, and the share price of Marine Harvest on the Oslo Stock Exchange. Considerable attention is paid to defining a relevant objective function for a hedger in the salmon market, and a new measure of hedging effectiveness is proposed. Among the candidate hedging instruments and their combinations, the salmon futures contract offers the highest hedging effectiveness; however, low liquidity may limit its applicability in practice.

In conclusion, the high uncertainty in the future spot price of salmon has been a constant predicament to the market participants and asks for a practical response. The research results contained in this thesis indicate that attempts of predicting the spot price might not deliver satisfactory results. However, hedging the price risk with salmon futures offers a substantial reduction in uncertainty and could therefore be promoted, provided that the futures contract attracts enough liquidity to meet the demand for hedging. The data used in the thesis is publicly and freely available, and the models are documented in detail; hence, they may be readily employed by the market participants in their business planning and optimization.

Contents

List of papers

This thesis is based on the following papers:

Paper 1

Salmon price volatility: a weight-class-specific multivariate approach

(Daumantas Bloznelis)

Paper 2

Short-term salmon price forecasting

(Daumantas Bloznelis)

Paper 3

Hedging salmon price risk

(Daumantas Bloznelis)

Acronyms

AIC	An information criterion, or Akaike's information criterion
AICc	Corrected AIC
ANN	Artificial neural network
AR	Autoregression, autoregressive
ARARMA	Autoregressive-autoregressive moving average
ARCH	Autoregressive conditional heteroskedasticity
ARFIMA	Autoregressive fractionally-integrated moving average
ARIMA	Autoregressive integrated moving average
ARMA	Autoregressive moving average
ASA	Allmennaksjeselskap, public stock-based company
BEKK	Baba, Engle, Kraft and Kroner (a type of multivariate GARCH model)
BIC	Bayesian information criterion, or Schwarz criterion
CBOT	Chicago Board of Trade
CEQ	Oslo Stock Exchange ticker for Cermaq ASA
DCC	Dynamic conditional correlation
DF-GLS	Modified Dickey-Fuller [test]
E.C.	European Community
EU	European Union
EUR	Euro
FAO	Food and Agriculture Organization of the United Nations
FGLS	Feasible generalized least squares
FPI	Fish Pool Index
FPSA15	Fish Pool Salmon Index 2015
GARCH	Generalized autoregressive conditional heteroskedasticity
GDP	Gross domestic product
GLS	Generalized least squares
GSF	Oslo Stock Exchange ticker for Grieg Seafood ASA
HAC	Heteroskedasticity-and-autocorrelation [robust standard errors]
ICSS	Iterated cumulative sums of squares
IPO	Initial public offering
ISA	Infectious salmon anemia
kNN, KNN	k-nearest neighbours

LARS	Least angle regression
LASSO	Least absolute shrinkage and selection operator
LSG	Oslo Stock Exchange ticker for Lerøy Seafood Group ASA
MA	Moving average
MAB	Maximum allowable biomass
MAE	Mean absolute error
MAFE	Mean absolute forecast error
MAPE	Mean absolute percentage error
MASE	Mean absolute scaled error
MHG	Oslo Stock Exchange ticker for Marine Harvest ASA
MT	Metric ton
NMBU	Norwegian University of Life Sciences
NOK	Norwegian krone
NOS	Norsk Opsjonssentral AS, a Norwegian financial service company
OBX	OBX Index, a stock market index consisting of the 25 most traded securities on the Oslo Stock Exchange, based on six months turnover rating
OHR	Optimal hedge ratio
OLS	Ordinary least squares
OSEBX	Oslo Børs Benchmark Index, an investible index which comprises the most traded shares listed on the Oslo Stock Exchange
RAM	Random-access memory
RMSE	Root mean squared error
RMSFE	Root mean squared forecast error
RRMSFE	Relative reduction in root mean squared error
SALM	Oslo Stock Exchange ticker for SalMar ASA
SSB	Statistics Norway
STL	Seasonal and Trend decomposition using Loess
TDN	A Norwegian news agency
TVINN	Norwegian Customs' electronic information system for the exchange of customs
U.K.	United Kingdom
U.S.	United States [of America]
VAR	Vector autoregression, vector autoregressive
VAR-EN	Vector autoregression estimated using elastic net regularization
VARIMA	Vector autoregressive integrated moving average

VARMA	Vector autoregressive moving average
VEC	Vector error correction
VECM	Vector error correction model
VECM-EN	Vector error correction model estimated using elastic net regularization
XEU	ISO code for the European Currency Unit (ECU), the predecessor of the euro

Introduction

The volatile price of salmon:
implications, responses and econometric modelling

1 Introduction

If the salmon spot price this week is NOK 40/kg, what will it be next week? How large a price change may one expect between two consecutive weeks? Should a salmon farmer rush to slaughter and sell the fish now or should he/she wait for a higher price? Given a futures price of NOK 35/kg for the next month, should a salmon processor buy some futures contracts if he/she is to reduce the uncertainty in the future costs of buying salmon?

Dealing with the uncertainty in the future spot price of salmon is the bread-and-butter activity of salmon farmers, processors, exporters, retailers and other market participants. An ability to accurately predict the price would simplify the management of a salmon farm or a processing plant, help optimize the operations and consequently increase the profit. Greater certainty over revenues from selling salmon in the future would likely make bank loans cheaper and increase the share prices of salmon farming companies. However, accurate price forecasts are hard to obtain; large price swings often take even the knowledgeable salmon experts by surprise (Grindheim, 2009, Jensen, 2011, Nilsen, 2011, Grindheim, 2013). It is also important to quantify the amount of uncertainty so that business managers may get prepared for the unknown of the future. At the same time, financial management tools such as forward contracting or hedging could be used to reduce the uncertainty in the futures spot price of salmon.

The aim of this thesis is to provide tools for answering the questions above and to illustrate the tools' use by employing them on historical data. The thesis consists of three single-authored papers, "Salmon price volatility: a weight-class-specific multivariate approach", "Short-term salmon price forecasting" and "Hedging salmon price risk". They take practical approaches towards solving the real-world problems of the salmon market participants. The ideas proposed in the papers could be implemented by a skilled econometrician or statistician for use with real-time data at a salmon farming, processing, export or retail company. They will hopefully yield greater efficiency in planning and management of the daily operations and help reduce the business financing costs.

The rest of the introductory chapter presents (2) the background of the research problems, (3) research problems specified in greater detail, (4) data and (5) methods used in the three papers, (6) summary of the three papers, and (7) contributions and limitations of the papers.

2 Background

This chapter provides an overview of the industry of and the market for the Norwegian farmed Atlantic salmon. The information is heavily based on Marine Harvest (2015) and Asche and Bjorndal (2011). These sources also contain more detailed content on the topics presented below as well as other topics related to salmon farming and the salmon market.

2.1 Historical overview of salmon farming

The origins of salmon farming stem from 1960s when experimental salmon farms were first set up in Norway (Marine Harvest, 2015). In the following decade farming was commercialized (Liu et al., 2011) and in 1980s the industry became commercially viable (Asche and Bjorndal, 2011, preface). The expansion was continuous and rapid; already by late 1990s farming overtook the wild capture in terms of production volume. Since 2012 the global production of farmed salmon has exceeded two million metric tonnes[1]; see Figure 1. With capture volume stagnating but farming volume continuously increasing, the share of farmed salmon in the total production has now reached about 75%.

Figure 1 Historical farmed salmon production volume, Norwegian (dark fill) and global, 1980-2013 (million tonnes)

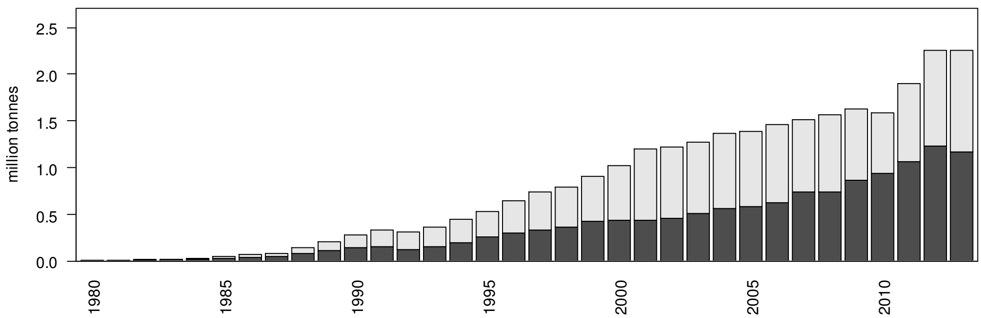

While the first commercially available farmed salmon was several times more expensive than its wild counterpart, continuous technological progress has allowed reducing the farmed salmon price over 1980s through early 2000s. Since then the price has been slowly increasing likely due to slower productivity growth, less rapid supply expansion (due to limited availability of production sites and

[1] The data is taken from Food and Agriculture Organization of the United Nations and can be accessed at http://www.fao.org/figis/servlet/SQServlet?file=/work/FIGIS/prod/webapps/figis/temp/hqp_5437944154593549544.xml&outtype=html

Figure 2 Norwegian farmed Atlantic salmon price, 1980-2013 (NOK/kg)

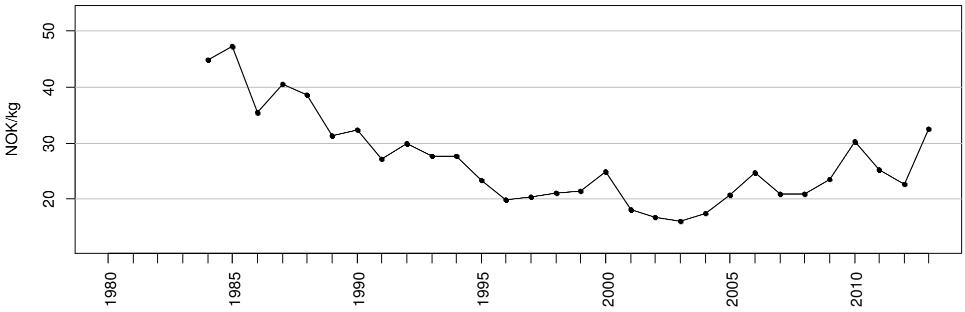

stricter industry regulation), increasing feed costs and soaring demand. The yearly Norwegian farmed salmon prices are depicted in Figure 2[2].

The salmon farming sector is an important element of the Norwegian industry. It provides jobs to some 22 700 people (Norwegian Seafood Federation and Norwegian Seafood Council, 2011), many of them in coastal communities with otherwise scarce employment opportunities. Salmon farming alone yields revenues in the order of NOK 40bn, which corresponds to between 1-2% of the Norwegian GDP. Export of farmed salmon from Norway is second only to the export of oil and gas.

Salmon is also produced in Chile, Scotland, Canada and to a smaller extent also other countries. The global production volume distribution as of 2014 is shown in Figure 3. As long as land-based farming remains only experimental, the natural conditions severely limit the geographical distribution of salmon producers in the world.

Figure 3 Distribution of farmed salmon production volume across countries, 2014

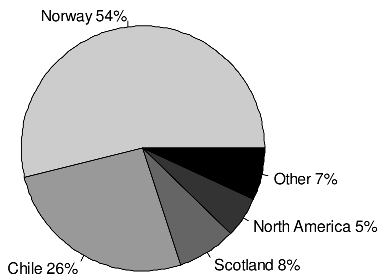

[2] The data is taken from Food and Agriculture Organization of the United Nations at http://www.fao.org/figis/servlet/SQServlet?file=/work/FIGIS/prod/webapps/figis/temp/hqp_5437944154593549544.xml&outtype=html and from Norges Bank at http://www.norges-bank.no/Statistikk/Valutakurser/valuta/USD/.

The salmon farming industry was characterized by small-scale operations in its early years, followed by gradual consolidation; the farming companies have merged repeatedly. Also, vertical integration has increased, leading to emergence of large companies that produce their own salmon feed, breed the salmon, farm, process and market it by themselves. The largest salmon farming companies are Marine Harvest, SalMar, Lerøy Seafood and Cermaq, all headquartered in Norway; Marine Harvest has farming divisions in all four major salmon farming countries, Cermaq in three (not in Scotland) while Lerøy Seafood and SalMar mainly operate in Norway. The largest five salmon farming companies have 58% of the total market share in Norway, 77% in Chile and over 90% in Canada and Scotland (Marine Harvest, 2015).

2.2 Salmon production cycle

The production cycle of farmed salmon lasts for about 24-36 months and is comprised of the following stages. First, eggs obtained from specially bred fish are cultivated in fresh water for 10-16 months until they become small fish weighing 60-100 grams. Second, the small fish is transferred to large nets in the sea where it is fed and continues to grow for another 14-22 months until reaching the weight of 3-6 kg. Third, the fish is harvested, slaughtered, gutted and packed; it might then either be further processed or sold to the end customer. For detailed schemes of the process, see Marine Harvest (2015).

From the production cycle it is clear that running a salmon farm requires relatively large initial investments and long-term planning; the final product becomes available only in two or three years following the decision to produce. The long production cycle and the resulting considerable delay in the reaction of supply to market events are likely causing low-frequency fluctuations in prices with prolonged periods of high prices followed by periods of low prices, and vice versa (Andersen et al., 2008).

2.3 Supply and demand for salmon

The supply of farmed salmon is determined by the availability of the production sites and industry regulation (very-long-term factors), amount of young fish put into the seawater for growing, the pace of growth and mortality (medium-term factors) and the farmer’s slaughtering schedule (medium- and short-term factor). Fish growth is conditional on water temperature (and therefore is seasonal), feeding intensity and schedule, and feed properties (Marine Harvest, 2015). While the seawater temperature is beyond the farmer’s control, the farmer can slow down or speed up growth by applying different feeding practices. Vaccination and other disease precautions as well as treatments against parasites are applied at different stages of production.

The main production risks consist in deterioration of biological conditions (e.g. toxic algae blooms); outbreaks of infectious disease (e.g. infectious salmon anemia decimated the Chilean salmon production in the late 2000s); unexpected development of water temperature; and winter storms that cause short-term disruptions in supply (Marine Harvest, 2015). Occasionally, the storms also damage the farming facilities in the sea leading to fish escapes; some of the fish may be recaptured but some are lost.

Naturally, the supply of salmon is also determined by its price. Farmers have short-term flexibility to respond to price changes within the limits of grown-up fish available for slaughtering at the time. As noted above, there is also a lagged effect of the currently expected future price on supply, the lag being due to the long production cycle. Andersen et al. (2008) and Asheim et al. (2011) are recent studies analyzing the short- and long-term effect of price on the salmon supply. A broad overview of salmon supply is available at Asche and Bjorndal (2011, p. 17-42).

The following factors affect the demand for salmon. First, the demand depends on the current price. Lagged price effects also are plausible demand factors as consumers may be persistent in their shopping habits and thus react to a change in price with a delay. Demand also depends on the price of substitutes, the closest substitute for farmed salmon being wild salmon (Asche et al., 1999, Asche et al., 2005) and trout (Asche and Bjorndal, 2011, p. 131-133); Asche and Bjorndal (2011, p. 131-133) note that salmon does not have any other close substitutes. Another demand factor is consumers' income (Asche and Bjorndal, 2011, p. 131-133).

Politics also plays a role in shifting the demand for salmon. Partial or complete trade restrictions for Norwegian salmon have been introduced and subsequently revoked repeatedly since the early years of salmon farming; for a detailed overview, see Asche and Bjorndal (2011, p. 139-141). Two prominent recent examples of trade restrictions are China's effective ban on the Norwegian salmon after the Nobel Peace Prize was awarded to a Chinese dissident Liu Xiaobo in 2010 and Russia's import ban on the Norwegian salmon in August 2014.

Media coverage may be affecting the demand for salmon by altering consumer tastes. There have been recurrent attacks on pollution and spread of parasites from salmon farms and lamentation of escapes of fish from the farms into the wilderness (Owen, 2008, Bloomberg, 2010, Coastal Alliance for Aquaculture Reform, Farmed Salmon Boycott, 2015). On the other hand, salmon has long been recognized as healthy food by dieticians (Lewin, 2014, Ipatenco, Boehlke, 2015), which should be a demand booster.

Both global and country-level demand for salmon has been analyzed in a number of studies, e.g. Bird (1986), Bjørndal et al. (1994), Asche (1996), Asche et al. (1998), Xie et al. (2009) and Xie and Myrland (2011).

2.4 Salmon trade

Salmon is a highly traded commodity; most of the salmon production is exported as the largest producer countries are different from the largest consumer countries. The map in Figure 4 (borrowed from Marine Harvest (2015)) shows salmon production, consumption and trade flows as of 2014. The Norwegian salmon mostly serves the European market with over two thirds of the Norwegian salmon being exported to the EU; 68% of the total Norwegian salmon export volume over 2010-2014 went to the EU.

The Norwegian salmon is traded in the spot market and using forward contracts. The exact volume shares of the two means of trade are unknown, but both are non-negligible. There is also a futures market for salmon, but no physical salmon is traded there as the futures contracts are cash settled.

Figure 4 Salmon production, consumption and trade flows in 2014. Source: Marine Harvest (2015)

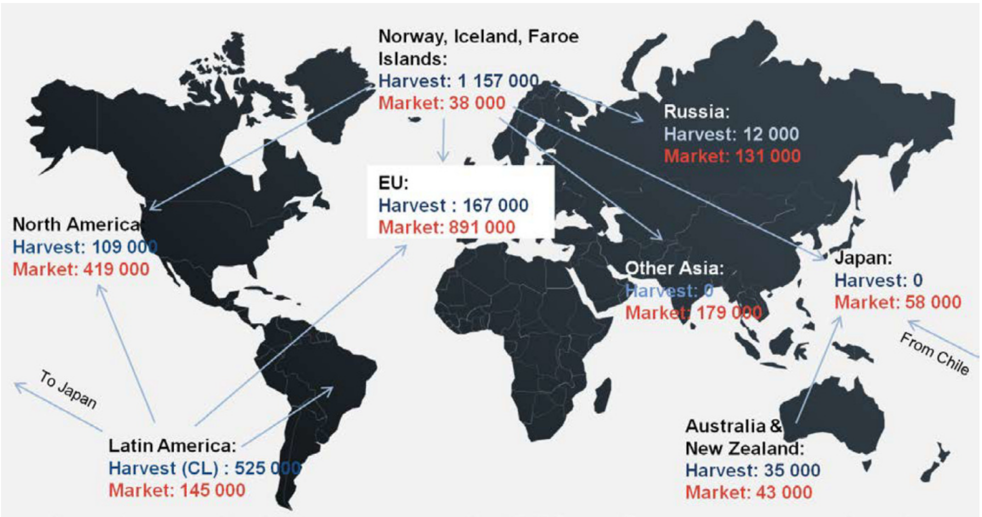

2.5 Financial derivatives on salmon price

Fish Pool futures and options exchange for salmon was established in mid-2006 in Bergen, Norway. It has since facilitated trade in financial derivates associated with the price of Norwegian farmed Atlantic salmon. The volume of trade in the futures contracts has expanded through 2006-2011 but stagnated and even fell in the subsequent years (Fish Pool, 2015b). Meanwhile, option trading seems to have had only negligible activity; however, public information on the trade volume for options is not available.

There exist futures contracts for each month up to 60 months ahead. It should be noted that the contracts with maturity longer than twelve months are traded very seldom; actually, even the nearby contracts have low trading volume. One may conjecture that the low liquidity of the market is partly due to the broad variety of available maturities so that there is a lack of concentration, which in turn repels potential buyers and sellers.

The contracts are settled against a monthly benchmark index called the Fish Pool Index (until the end of 2014) or the FPSA15 (from the beginning of 2015) (Fish Pool, 2015c). The monthly index is an equally-weighted average of the corresponding weekly index values; each trading month has either four or five complete weeks, as detailed in the trading rules (Fish Pool, 2015a). The weekly index is composed of a few price indices that should reflect the spot price of 3-6 kg salmon. Historically, the weights of the different component indices have varied, but the dominating component has always been the spot price by NOS, later replaced by NASDAQ (see chapter 4 for details). The current composition of the weekly index can be found at the Fish Pool website (Fish Pool, 2015c).

Price quotes from the futures market have been studied by Solibakke (2012), Ewald (2013), Ewald et al. (2014) and Ewald and Salehi (2015). An overview of the perspectives of a futures market for salmon is given in Bergfjord (2007).

3 Research problems and objectives

As indicated in the chapter 1, knowledge of the future spot price is important for business planning for all the salmon market participants: farmers, processors, exporters, and retailers. Since a considerable share of salmon is traded in the spot market, the spot price is a key determinant of a salmon farmer's revenues and a processor's, exporter's or retailer's costs. The spot price is considerably more volatile and more difficult to predict than production or consumption volume. Hence, it is the spot price that constitutes the lion's share of uncertainty in the future revenues and costs of the salmon market participants. From a salmon farmer's perspective, the fluctuations in the spot price bring about not only quantitative but also qualitative effects to the financial accounts; they may determine whether the net result will be a profit or a loss. This is even more true for salmon processors; while salmon farmers have enjoyed prolonged periods with positive net results, salmon processors have had numerous periods of losses. In sum, developing a deeper understanding of price volatility and better forecasting models for the future spot price of salmon is vital.

There are two recent studies of salmon price volatility, Oglend and Sikveland (2008) and Oglend (2013), who document the properties of the volatility of the average spot price of salmon. Oglend (2013) notices and discusses a structural change in the volatility around 2006 but fails to properly account for it in the econometric models. Neither of the studies considers the volatilities of the spot prices of specific weight classes. Both of these issues are taken into consideration in Bloznelis (2016b) which aims to detect and describe patterns in the salmon price volatility and compare them across the weight classes.

Meanwhile, the literature on salmon price forecasting consists of Lin et al. (1989), Vukina and Anderson (1994), Gu and Anderson (1995), and Guttormsen (1999). Only the latter study focuses on the spot price of the Norwegian salmon. With the forecasting studies dating back to the previous millenium, the time is ripe for an update, which I provide in Bloznelis (2016c). The paper aims for carrying out a realistic forecasting exercise, assessing forecast accuracy and optimality, and drawing conclusions and implications for the market participants.

Alternatively, the uncertainty in the future spot price of salmon could be mitigated through hedging; that is, investing in financial instruments whose prices move in a direction opposite to that of the salmon price. The purpose is to construct a portfolio of physical salmon and a hedging instrument such that the value of the portfolio would fluctuate less, and be more predictable than, that of physical salmon. To my knowledge, this topic has not been investigated before in the academic literature. Bloznelis (2016a) is intended to fill this gap; it is an exploratory study of hedging the salmon price risk with various hedging instruments using different hedging strategies. Finding relevant hedging instruments and optimal hedge ratios, assessing hedging effectiveness, and comparing different hedging strategies are the main objectives of the paper.

4 Data

4.1 Sources

Data sources relevant to modelling the salmon price and its volatility are rich and generally of high quality. The data series and their sources are presented in Table 1. Columns “Range” and “Frequency” indicate the specific ranges and frequencies used in the three research papers, but not necessarily all the ranges and frequencies available at the sources. All of the data used in this thesis is publicly and freely available. I used weekly data is my models; hence, most of the data taken from the sources is in weekly frequency. When weekly data was not available, I took daily data and converted it into weekly before modelling.

4.2 Problems with the data

Most of the data is clean and orderly. However, I have encountered the following few issues. First, in the NOS dataset there are three corrupt price values (one for the price of the 1-2 kg fish and two for 7+ kg fish). When using this data in Bloznelis (2016b), I imputed values equal to the equally-weighted average of the preceding and the succeeding values. This imputation method can be justified for cases where only a very small proportion of values is missing; in my applications, the

Table 1 Data sources

Object	Range	Frequency	Source	Web address
Salmon spot price (NOS price)	1995w1 – 2013w13	Weekly	NOS	https://salmonprice. nasdaqomxtrader.com/ public/home
Salmon spot price (NASDAQ price)	2013w14 – 2015w17	Weekly	NASDAQ	https://salmonprice. nasdaqomxtrader.com/ public/home
Salmon export price (SSB price)	2007w1 – 2014w39	Weekly	Statistics Norway (SSB)	www.ssb.no
Salmon export volume (SSB volume)	2007w1 – 2014w39	Weekly	Statistics Norway (SSB)	www.ssb.no
Salmon futures prices	2007-01-01 – 2015-04-26	Daily	Fish Pool	www.fishpool.eu
Fish Pool Index (FPI, FPSA15)	2007-01-01 – 2015-04-26	Weekly	Fish Pool	www.fishpool.eu
Salmon production volume (global and Norwegian)	1980 – 2013	Yearly	FAO	www.fao.org
Salmon production value (global and Norwegian)	1980 – 2013	Yearly	FAO	www.fao.org
XEU/NOK exchange rate*	1995-01-01 – 1998-12-31	Daily	Norges Bank	www.norgesbank.no
EUR/NOK exchange rate	1999w1 – 2014w39	Weekly	Oanda	www.oanda.com

Note: "w" denotes week, e.g. 1995w1 denotes 1995 week 1. Ranges and frequencies as used in the three research papers; wider ranges and/or extra frequencies may also be available at the sources.

*XEU is the ISO code for the European Currency Unit (ECU), the predecessor of the euro.

proportions were one and two observations out of around five hundred. Second, in the daily futures price data from Fish Pool, there are a few cases of observations not following the chronological order; care needs to be taken when employing automated routines that would normally assume properly ordered inputs. Third, an outlier observation indicating an unusually low value of the FPI on 2010 week 52 can be found in the Fish Pool dataset. It is a technical mistake that has been introduced only recently, while the mistake-free version of the FPI series used in Bloznelis (2016b) and Bloznelis (2016a) was still available at the source in mid-2013 when the data for the studies was retrieved.

In empirical studies such as this thesis, good understanding of how the data was generated is crucial for its appropriate use. Therefore, considerable attention has been paid to explaining the nature of the data in all three papers. Below I present a summary of the relevant material. Starting with the spot prices, the best available spot price series is perhaps the NOS series for 1995 week 1 through 2013 week 13 and the NASDAQ series starting in 2013 week 14. The NOS and NASDAQ prices are survey prices and as such do not cover the entire spot market. The NOS survey used to cover around one third of the production volume and the same is true for the NASDAQ survey. Also, no explicit mechanism for ensuring fair reporting appears to exist; however, this does not seem to be of major concern in practice.

There is a certain mismatch of definitions of the NOS and the NASDAQ prices. While NOS reflects prices paid by exporters to salmon farmers, NASDAQ indicates prices received by salmon exporters from foreign buyers. Thus the two prices differ by the exporter's margin which has been estimated at NOK 0.75/kg as of late 2012 – early 2013 (Fish Pool, 2014). Also, the conditions for inclusion in the NASDAQ survey are less stringent than those for the NOS survey; hence, NASDAQ should cover a larger spot market share than NOS; however, this does not hold in practice. A more detailed description of the two series can be found in Bloznelis (2016c) and Fish Pool (2014). Although it is common to concatenate the two series adjusting for the NOK 0.75/kg margin, the intrinsic differences between them should be kept in mind. For example, Bloznelis (2016c) notes that using the data up to 2013 week 13 to estimate forecasting models for predicting beyond 2013 week 13 may lead to reduced forecasting accuracy.

With regards to the FPI/FPSA15, its composition has changed multiple times since the creation of the index. Some component series have been assigned higher or lower weights (NOS and NASDAQ prices, SSB price) while others which initially were included have later been completely eliminated (Farmers index, Mercabarna index, Rungis index). If explanatory modelling of the FPI was to be undertaken, changes in its composition should be accounted for.

Daily futures prices are rather ephemeral due to the low liquidity of the futures market, as discussed in chapter 2. This is little relevant for Bloznelis (2016c) where the existence of a price quote is sufficient; meanwhile, it is quite important for Bloznelis (2016a) where the ability to trade at the given price is crucial. The liquidity problem seems to have been overlooked in essentially all major studies on salmon futures published before 2016, e.g. Ewald (2013), Ewald et al. (2014), Ewald and Salehi (2015), and Solibakke (2012).

Finally, some of the data reported in Table 1 is no longer available at the sources, e.g. the weekly XEU/NOK rate.

5 Methods

5.1 To explain or to predict?

Econometric modelling strategies can be categorized by their purpose into (1) explanatory (E) and (2) predictive (P). Explanatory models are intended for causal explanation of processes or relationships under examination, while predictive models are used for empirical forecasting. (A third category – descriptive modelling – could be distinguished as well; however, it is less popular in economics and therefore is omitted from the following discussion.) The division between (E) and (P) can be briefly summarized by answering the question, To explain or to predict?, raised in Shmueli (2010). There are a number of differences between (E) and (P). First, the focus of (E) is analysis of causation while the focus of (P) is analysis of association. Second, functional relationships used in (E) are theory driven and support explanation and causal hypotheses testing; meanwhile, in (P) they are data driven and need not be directly interpretable. Third, modelling is backward looking in (E) as it involves testing a predetermined set of hypotheses; on the contrary, modelling is forward looking in (P) with the goal of predicting new observations. For a more formal treatment, let us assume a true but unknown underlying functional relationship f between an outcome y and inputs x, $y = f(x)$. When modelling the relationship, a function $\hat{f}$ is used as an approximation of f. Given a quadratic loss function, the expected prediction error, EPE, for a new observation of inputs is

$$
\begin{aligned}
EPE &= E\left(y - \hat{f}(x)\right)^2 \\
&= E\big(y - f(x)\big)^2 + \left(E\left(\hat{f}(x) - f(x)\right)\right)^2 + E\left(\hat{f}(x) - E\left(\hat{f}(x)\right)\right)^2 \\
&= Var(y) \qquad + \left(Bias\left(\hat{f}(x)\right)\right)^2 \qquad + Var\left(\hat{f}(x)\right) \\
&= \sigma^2 \qquad\quad\;\; + \left(Bias\left(\hat{f}(x)\right)\right)^2 \qquad + Var\left(\hat{f}(x)\right).
\end{aligned} \tag{1}
$$

See Hyndman (2015) for a proof. The goal of (E) is to minimize the squared bias term $\left(Bias\left(\hat{f}(x)\right)\right)^2$ and ideally to get an unbiased estimator $\hat{f}(x)$ of $f(x)$. Note that (E) ignores the estimation variance $Var\left(\hat{f}(x)\right)$ which reflects how well we can estimate $\hat{f}$ given its general form, the unknown parameters and a data sample. The estimation variance arises from our inability to have perfect estimation precision given a data sample of limited size. Meanwhile, the goal of (P) is to minimize the expected prediction error. That amounts to joint minimization of the squared bias and the variance, while σ^2 is a constant and hence can be excluded from consideration. Since there is a trade-off between the squared bias and the variance, the joint minimization due to (P) may (and generally will) yield higher squared bias than obtained in (E). That is, (P) may require an increase in the squared bias relative to (E) if that yields a reduction in the expected prediction error. A more detailed exposition of the bias-variance trade-off may be found in Hastie et al. (2009, p. 223-228). Shmueli (2010) offers a more complete comparison of the explanatory and the predictive modelling approaches. For an alternative perspective on methodological division of statistical modelling see Breiman (2001).

Econometric modelling in economics has almost exclusively been explanatory rather than predictive (Shmueli, 2010). Explanatory modelling may well suit the purpose of enhancing the general understanding of processes of interest and/or drawing policy implications. However, should forecasting be considered less important than that? What can be more rewarding and practical than knowing, or having a good insight into, the future? The problem of the uncertainty in the future spot price of salmon could be viewed from both the explanatory and the predictive perspective. First, regulators and market participants wondering why the price volatility is high and what the mechanism bringing about the large price swings is would have to resort to explanatory modelling. This has been done, at least partly, in Oglend (2013). Meanwhile, Bloznelis (2016b) offers a more detailed descriptive picture of the subject. Second, from the viewpoint of a market player without the power to change the price formation process, the problem of high price volatility naturally calls for a predictive approach. The response is the forecasting exercise in Bloznelis (2016c) and the out-of-sample hedging study in Bloznelis (2016a). Hence, the three papers in this thesis could be classified as descriptive, predictive and predictive, respectively.

5.2 Modelling time series for short-term forecasting

The data used in this thesis is of time-series type. Two strands of time series modelling may be distinguished, (1) structural modelling and (2) reduced-form modelling. Structural modelling is the time-series-specific counterpart of explanatory modelling. Theoretically-grounded structural models are used for inference with respect to the structural parameters that directly correspond to theoretical quantities. Meanwhile, reduced-form models are not directly suited to statistical inference and are either converted to structural models in case of explanatory modelling or used directly for prediction. They might also be used in a descriptive context as in Bloznelis (2016b). Another difference between the structural and the reduced-form time series models is that the former may perform well in long-term forecasting where certain equilibrium or structural relationships prevail; meanwhile, the latter may dominate in the short term where these relationships are overshadowed by random shocks. The short-term focus in the three papers of this dissertation warranted the use of the reduced-form models. There are also black-box techniques such as neural networks and the k-nearest neighbour method used in Bloznelis (2016c). These may be characterized as algorithmic modelling approaches while the time series models mentioned above belong to the data modelling approach, as defined in Breiman (2001).

5.3 Main models

I will overview a few cornerstone models that are encountered repeatedly in this thesis, leaving out a number of models used only in Bloznelis (2016c). These models can be classified into conditional mean and conditional variance models.

Regarding the conditional mean models, regression with autoregressive moving-average (ARMA) errors is used for seasonal adjustment of the price data in all three papers; autoregressive integrated moving-average (ARIMA) model is used in Bloznelis (2016b) and Bloznelis (2016c) to model the conditional mean of the spot price of salmon; vector error correction model (VECM) is used in Bloznelis (2016c) and Bloznelis (2016a) to model the conditional mean of a system of cointegrated variables, one of which is the average spot price of salmon.

5.3.1 ARIMA

The ARIMA model is a classic and perhaps the most popular of all time series models. It was introduced in Whittle (1951) and popularized by Box and Jenkins (1970). An overview of ARIMA models can be found in any time series textbook. A classic treatment is Box et al. (2011), also Hamilton (1994).

Consider a model of the form

$$x_t = \varphi_0 + \varphi_1 x_{t-1} + \cdots + \varphi_p x_{t-p} + \varepsilon_t + \theta_1 \varepsilon_{t-1} + \cdots + \theta_q \varepsilon_{t-q}. \tag{2}$$

The first p terms after the intercept are lags of the dependent variable, or autoregressive terms, and the model's autoregressive order is said to be p. The last q terms are lags of the error (or shock, or innovation), or moving-average terms, where q is the model's moving-average order. The errors are independently and identically distributed and often assumed to have a normal distribution. A model of such form is shortly denoted ARMA(p,q). If $x_t = y_t - y_{t-1} =: \Delta y_t$ and x_t follows an ARMA(p,q) model, then y_t is said to follow an ARIMA (p,d,q) model with an integration order $d = 1$. Higher orders of integration $d > 1$ are rarely encountered in the economics literature and were not used in the papers of this thesis.

Let us consider two simple special cases that should convey the intuition behind the model. An ARIMA(1,0,0) model without an intercept is $x_t = \varphi_1 x_{t-1} + \varepsilon_t$. The value of x at time t is just a multiple of the preceding value of x, plus a random shock. The autoregressive coefficient φ_1 in such a model normally satisfies the condition $|\varphi_1| < 1$. Thus the process $\{x_t\}$ tends to return gradually to zero if not for the random shocks. The trajectory of the process could be thought of as a discretized path of a cyclist following a straight road with a number of random-sized potholes right in the middle; the ultimate direction is straight ahead, but avoiding potholes requires temporarily but repeatedly missing the straight line. Another simple example is an ARIMA(0,0,1) model without an intercept, $x_t = \varepsilon_t + \theta_1 \varepsilon_{t-1}$; normally $|\theta_1| < 1$. The model reflects delayed reactions to shocks. x_t

would be zero if not for contemporaneous as well as lagged innovations. When a shock of size ε_t hits at time t, its effect is felt at time $t+1$ as well, but is scaled. That may happen when a market underreacts or overreacts to a shock at the time it hits but corrects itself towards a more "reasonable" reaction in the next period. For example, if news reach some of the market participants immediately but others with a one-period delay, the observed market reaction could be characterized by an ARIMA(0,0,1) process.

ARIMA models have the advantage of parsimony; however, their coefficients are not directly informative unless in the very simples cases such as ARIMA(p,0,0) or ARIMA(0,0,q). Two ARIMA models with quite different autoregressive and moving average orders and dissimilar coefficient values may generate processes with very similar characteristics. The behaviour of y_t following an ARIMA(p,d,q) process is best reflected by impulse-response functions and their graphs rather than model coefficients.

The model can be estimated via the maximum likelihood method which gives precise estimates but is computationally intensive. There are compromise methods such as conditional maximum likelihood that sacrifice some precision (but not consistency) to gain estimation speed. Conditional maximum likelihood estimation was utilized in the papers contained in this thesis.

In Bloznelis (2016b), ARMA models were used as the conditional mean models for logarithmic returns on the salmon spot prices for different weight classes. In Bloznelis (2016c), ARIMA models were used to model the average spot price.

5.3.2 Regression with ARMA errors

The regression with ARMA errors is represented by the following two equation system

$$y_t = \beta_0 + \beta_1 x_{1,t} + \cdots + \beta_K x_{K,t} + u_t \tag{3a}$$

$$u_t = \varphi_1 u_{t-1} + \cdots + \varphi_p u_{t-p} + \varepsilon_t + \theta_1 \varepsilon_{t-1} + \cdots + \theta_q \varepsilon_{t-q} \tag{3b}$$

where y is the dependent variable, x_1 through x_K are independent variables and ε is independently and identically distributed error. When the model is used for seasonal adjustment in Bloznelis (2016b), Bloznelis (2016c) and Bloznelis (2016a), the x's in equation (3a) are Fourier terms and dummy variables representing the seasonal effects. In particular, up to 26 pairs of $sin(\cdot)$ and $cos(\cdot)$ waves with increasing frequencies constitute the Fourier terms; four dummies are used for the weeks around Christmas (two weeks before, the Christmas week and one week after); and analogous four dummies account for the seasonal effect due to Easter. Allowing the model error u_t to have an ARMA(p,q) structure accommodates the possible short-term patterns in u_t. This approach to seasonal adjustment of weekly data has been recommended in Hyndman (2010b) and Hyndman (2014). The model is estimated similarly to the ARIMA model as discussed above. An overview of regression with ARMA errors and a few related models is given in Hyndman and Athanasopoulos (2014) and Hyndman (2010a).

5.3.3 VECM

The vector error correction model (VECM) is a multivariate, regression-based time-series model applicable in the context of cointegrated time series. Let me first introduce the concepts of integration and cointegration. Consider a pair of stationary time series z_τ and w_τ. Define a new pair of time series, x_t and y_t, as cumulative sums of z_τ and w_τ, respectively; i.e. $x_t := \sum_{\tau=0}^{t} z_\tau$ and $y_t =: \sum_{\tau=0}^{t} w_\tau$. The new series x_t and y_t will by construction be nonstationary and integrated, the latter term reflecting the nature of x_t and y_t as cumulative sums, or integrals. If a linear combination $x_t + \beta y_t$ of two integrated time series is stationary, the two series are said to be cointegrated. For example, x_t could stand for the salmon spot price and y_t for the salmon futures price. If their linear combination, $x_t - y_t$, were stationary, the spot price and the futures price would be cointegrated. The difference between the spot and the futures prices being a stationary process implies that the spot price tends to be close to the futures price, and vice versa. (Cointegration can be extended to include more than two variables and higher orders of integration, but it is unnecessary for this exposition.) Cointegrated variables are often modelled using the VECM. In a bivariate case, the VECM takes the following form:

$$\Delta \binom{x}{y_t} = \binom{\alpha_1}{\alpha_2} (1 \quad \beta) \binom{x_{t-1}}{y_{t-1}} + \Gamma_1 \Delta \binom{x_{t-1}}{y_{t-1}} + \cdots + \Gamma_\mathrm{p} \Delta \binom{x_{t-p}}{y_{t-p}} + \binom{u_t}{v_t}; \tag{4}$$

here Δ denotes difference operator such that $\Delta x_t := x_t - x_{t-1}$; α_1, α_2 are constants called loading coefficients; $(1 \quad \beta)$ forms the (transposed) cointegrating vector such that for a pair of integrated processes (x_t, y_t), $(1 \quad \beta) \binom{x_t}{y_t}$ is a stationary process; Γ_1 through Γ_p are 2×2 coefficient matrices; and u_t and v_t are error terms. If x_t and/or y_t contains a linear trend, the model will also include a constant term.

VECM can be estimated using either maximum likelihood or ordinary (or generalized) least squares techniques (OLS or GLS, respectively), or their combinations; see. e.g. Lütkepohl (2007, p. 269-322). A comprehensive overview of the model is given in Lütkepohl (2007, p. 237-384) and Johansen (1995, p. 70-209).

Conditional variance of multivariate time series is modelled in Bloznelis (2016b) and Bloznelis (2016a). Dynamic conditional correlation-generalized autoregressive conditional heteroskedasticity (DCC-GARCH) model was used in both cases. The model is comprised of two parts that will be discussed sequentially.

5.3.4 GARCH

Volatility clustering is one of the stylized facts of financial returns. Large swings in prices are typically followed by large swings, but small moves are succeeded by small (Cont, 2001). Engle (1982) proposed the following autoregressive conditional heteroskedastic (ARCH) model to accommodate these features:

$$u_t = \sigma_t \varepsilon_t \tag{5a}$$

$$\sigma_t^2 = \omega + \alpha_1 u_{t-1}^2 + \cdots + \alpha_m u_{t-m}^2 \tag{5b}$$

where $\varepsilon_t \sim i.i.d(0,1)$. It is a deterministic model in which the conditional variance of the error term at time t is completely determined by the magnitudes of the last m errors. The model was later generalized in Bollerslev (1986) who proposed the generalized ARCH (GARCH) model:

$$u_t = \sigma_t \varepsilon_t \tag{6a}$$

$$\sigma_t^2 = \omega + \alpha_1 u_{t-1}^2 + \cdots + \alpha_m u_{t-m}^2 + \beta_1 \sigma_{t-1}^2 + \cdots + \beta_s \sigma_{t-s}^2 \tag{6b}$$

where $\varepsilon_t \sim i.i.d(0,1)$, which is the regular ARCH model with added terms of lagged conditional variances. GARCH(1,1) with one lagged squared error and one lagged conditional variance is the most popular version of GARCH models and likely the most popular conditional variance model in general. It is found to fare well empirically against its competitors (Hansen and Lunde, 2005). There are many extensions and variations to the standard GARCH model, see e.g. Bollerslev (2009) or Teräsvirta (2009). Given a time series, the relevance of a GARCH model for the conditional variance can be assessed by testing for presence of autoregressive conditional heteroskedasticity, e.g. by using the ARCH-LM test (Engle, 1982). The model is estimated via the method of maximum likelihood.

5.3.5 DCC

While ARCH and GARCH models are directly suitable for modelling the conditional variance of univariate time series, generalizing them to the multivariate setting has proven to be challenging. A direct multivariate counterpart of the univariate GARCH model is the VECH-GARCH model (Bollerslev et al., 1988). VECH-GARCH suffers from high computational complexity due to its large number of parameters; in addition, it is problematic to ensure that the estimated conditional variance matrix is positive semidefinite. A number of alternatives to the VECH-GARCH model have been proposed; they are overviewed in Bauwens et al. (2006) and Silvennoinen and Teräsvirta (2009). Bauwens et al. (2006) classify the alternative multivariate GARCH models into (1) direct generalizations of the univariate GARCH model (e.g. VECH-GARCH, BEKK-GARCH and factor GARCH); (2) linear combinations of univariate GARCH models (e.g. orthogonal GARCH, generalized orthogonal GARCH and latent factor GARCH); and (3) nonlinear combinations of univariate GARCH models (e.g. dynamic conditional correlation GARCH, general dynamic covariance model and copula-GARCH). Silvennoinen and Teräsvirta (2009) provide an alternative classification into (1) models of the conditional variance matrix (e.g. VECH-GARCH and BEKK-GARCH); (2) factor models (e.g. factor GARCH and (generalized) orthogonal (factor) GARCH); (3) models of conditional variances and correlations (e.g. (extended) constant conditional correlation GARCH and dynamic conditional correlation GARCH); and (4) nonparametric and semiparametric models (e.g. semiparametric conditional correlation GARCH). None of the alternative multivariate GARCH models seems to clearly dominate the rest in terms of both flexibility and computational ease. Bloznelis (2016b) and Bloznelis (2016a) use dynamic conditional correlation GARCH (DCC-GARCH, or just DCC) model proposed in Engle (2002) and Engle and Sheppard (2001). The model choice is based on practical considerations of availability of the relevant statistical software, estimation speed and ease of interpretation.

Given a multivariate time series with a zero conditional mean vector, the DCC model starts with considering the conditional variance of each of the (univariate) component series. A univariate GARCH (or some other) model is estimated for each univariate series, and fitted conditional standard deviations are extracted. Each series is then standardized by dividing through the fitted conditional standard deviations. The correlations between pairs of the standardized series are then represented in the spirit of a GARCH model. Define conditional quasi-correlation at time t as

$$q_{i,j,t} = (1 - \gamma - \delta)\rho + \gamma\varepsilon_{\mathrm{i},t-1}\varepsilon_{\mathrm{j},t-1} + \delta q_{i,j,t-1} \tag{7}$$

with unconditional correlation ρ and lagged conditional quasi-correlations q_{t-1}. Then the proper conditional correlation is obtained as scaled conditional quasi-correlation:

$$\rho_{i,j,t} = \frac{q_{i,j,t}}{\sqrt{q_{i,i,t}}\sqrt{q_{i,i,t}}}. \tag{8}$$

Scaling ensures the conditional correlation lies strictly between negative one and one.

The model is estimated in two stages (first the univariate GARCH models and then the conditional correlation part) by the maximum likelihood method. An enlightening critique of the DCC-GARCH model has been given in Caporin and McAleer (2013).

5.4 Rolling windows for out-of-sample performance evaluations

The forecasting exercise in Bloznelis (2016c) and the hedging study in Bloznelis (2016a) call naturally for out-of-sample performance evaluations. Forecasts cannot be made using the data that is to be predicted, and minimizing uncertainty is a meaningful concept only when the future is unknown. To make this restriction operational, the original sample may be split into a training subsample to be used for model building followed by a validation (or holdout) subsample for measuring forecasting or hedging performance. This allows carrying out a realistic simulation of forecasting or hedging that mimics the information availability of the real world setting. However, a single sample split will only produce one forecast or one optimal hedge ratio for a particular forecast/hedging horizon. Performance evaluation will be more reliable and robust if more than one instance of out-of-sample performance is obtained. For that matter rolling window approach is used in Bloznelis (2016c) and Bloznelis (2016a). That is, the T-long original sample is divided into m overlapping subsamples that are $T - m + 1$ observations long, as illustrated in Figure 5 for $T = 10$, $m = 6$.

Consider forecasting the salmon price one period ahead. Each of the rolling windows (consisting of grey-shaded cells in Figure 5) can be used for estimating forecasting models, which in turn can produce one-period-ahead forecasts. The forecasts can be compared with the realized values (the black-shaded cells in Figure 5) for all the windows except for the last one. This way the use of the rolling windows allows assessing the forecast accuracy in a fair way; it helps avoid the use of data that is not available at the time of making the forecast. Similarly, hedging effectiveness can be assessed in a fair manner when based on an out-of-sample evaluation, just like in the case of forecasting. Note that the rolling window technique increases the computational costs slower than linearly in the number of windows m as compared with the benchmark cost of running the estimation once on the original sample.

As noted in Bloznelis (2016c), the use of rolling windows for model selection will systematically favour models that are too parsimonious for the original sample. Intuitively, a window is smaller than the entire sample, and there is less data available for parameter estimation. With less data, the estimation precision decreases and the associated variance of the estimated models (as discussed in the subchapter 5.1) increases. This tends to bias the model selection towards more parsimonious models that have lower variance than the richer models. In other words, reduced sample size yields increased estimation variance, and it makes it seem optimal to cut down the model complexity to trade off an increase in the squared bias for a reduction in variance. This flaw and any responses to it remains an open research problem.

The rolling window framework allows models to differ across the windows, which lends increased robustness against potential structural changes in the underlying data generating process. For an illustration, suppose there is a structural break in the middle of the sample; suppose also that the rolling window is shorter than half of the sample (see Figure 6). Then the first rolling window will not be affected by the data at and after the break, and therefore will reflect the homogenous behaviour of the process before the break. Similarly, the last rolling window will not include the relevant data from at or before the break and will only cover a homogenous period at the end of the sample. The windows in between may fall entirely into the pre- or post-break period or may cover the break and both part of the period before it and part of the period after it. However, at least the first and the last windows will avoid any detrimental effects from a neglected structural break.

Figure 5 Original sample and rolling windows within it

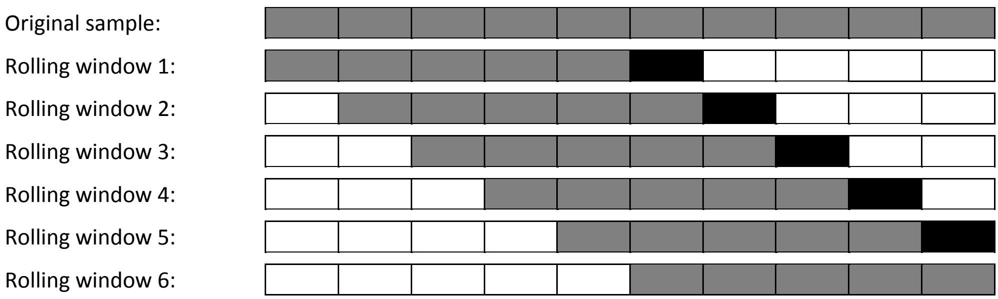

Note: grey cells mark inclusion into the training subsamples. Black cells mark the validation (or holdout) subsamples, i.e. the units suitable for out-of-sample forecast evaluation of one-step-ahead forecasts and hedging effectiveness evaluation of one-step-ahead hedges based on the training subsamples.

Figure 6 Original sample with a structural break in the middle, and rolling windows

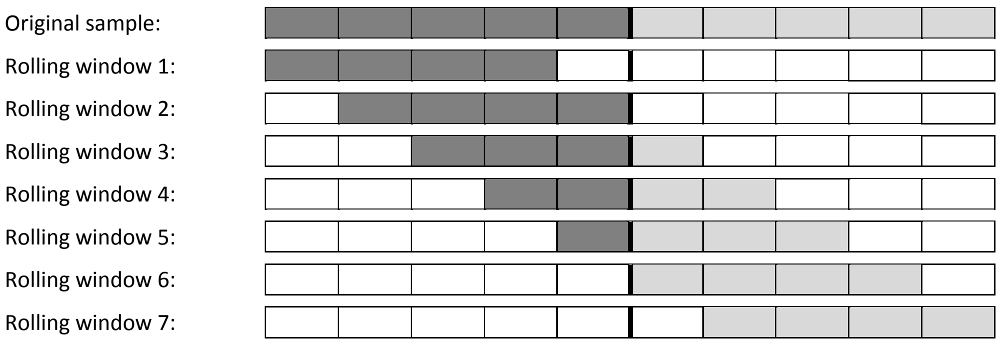

Note: dark grey cells mark the sample points due to the data generating process before the structural break. Light grey cells mark the sample points due to the data generating process after the structural break. While the original sample includes the pre- and post-break periods and thus the break needs to be accounted for in the modelling, the first and the last rolling windows are not affected by the structural break and will reflect the data generating processes before and after the break, respectively.

5.5 Variable selection and regularization

Model selection, and/or variable selection have been confronted repeatedly in Bloznelis (2016b), Bloznelis (2016c) and Bloznelis (2016a); it constituted a crucial modelling step in each case. Reduced-form time series modelling is characterized by a lack of theory-implied restrictions on the shape of the models, the lag orders and the sensitivity of the outcome variable to the input variables. However, the importance of selecting the relevant variables or lag orders is as great as ever; reliability of model-based inference (in explanatory modelling), characterization of the process (descriptive modelling) as well as forecast accuracy (predictive modelling) depend heavily on model and variable selection.

5.5.1 Selection and regularization techniques

The most common methods of variable selection in regression or regression-based models in economics, aside from theoretical argumentation, are significance testing, forward and backward stepwise selection and use of information criteria. Except for the latter one, these techniques have well-known flaws and have been largely abandoned in the modern statistics and machine learning literature. Taking a step towards contemporary methods one may refer to Hastie et al. (2009, p. 57-94) to find three categories of variable selection techniques for linear regression: subset selection, regularization (shrinkage) and methods using derived input directions. Among the three, subset selection is perhaps the most attractive for the purposes of explanatory modelling due to the explicit division of variables and/or lags into "relevant" and "irrelevant" ones. By a tautological definition, "relevant" variables/lags are the ones from the pool of all candidate variables/lags that happen to be included in the final model, and their inclusion is based on explanatory power; "irrelevant" variables/lags are the ones remaining excluded.

However, regularization may appear more natural than subset selection when it comes to predictive modelling. Regularization allows for greater flexibility in that it may include, exclude, or *partly* include a variable; it is thus a generalization of subset selection. Ridge regression (Hoerl, 1962, Hoerl and Kennard, 1970) is perhaps the oldest example of regularization. When estimating a linear regression model

$$y_i = \beta_0 + \beta_1 x_{1,i} + \cdots + \beta_K x_{K,i} + \varepsilon_i \tag{9}$$

via ridge regression, the following objective function is being minimized with respect to the parameters β_1 through β_K:

$$\sum_{i=1}^{N} \left(y_i - \beta_0 - \beta_1 x_{1,i} - \cdots - \beta_K x_{K,i}\right)^2 + \lambda \sum_{k=1}^{K} \beta_k^2 \tag{10}$$

where λ is a tuning parameter that is taken as given. The ridge estimator is the vector

$$\hat{\beta}^{ridge} = \mathrm{argmin}_{\beta} \left(\sum_{i=1}^{N} \left(y_i - \beta_0 - \beta_1 x_{1,i} - \cdots - \beta_K x_{K,i}\right)^2 + \lambda \sum_{k=1}^{K} \beta_k^2\right). \tag{11}$$

It is easy to see that ridge regression has the same objective function as the regular linear regression – but with a quadratic penalty on the size of the regression coefficients (except for the intercept). Least absolute shrinkage and selection operator (LASSO) (Tibshirani, 1996) is another popular regularization technique similar to ridge regression. The difference between the two is that LASSO applies a linear (rather than quadratic) penalty in the objective function,

$$\sum_{i=1}^{N} \left(y_i - \beta_0 - \beta_1 x_{1,i} - \cdots - \beta_K x_{K,i}\right)^2 + \lambda \sum_{k=1}^{K} |\beta_k|, \tag{12}$$

to yield the LASSO estimator

$$\hat{\beta}^{LASSO} = \underset{\beta}{\operatorname{argmin}} \left(\sum_{i=1}^{N} \left(y_i - \beta_0 - \beta_1 x_{1,i} - \cdots - \beta_K x_{K,i} \right)^2 + \lambda \sum_{k=1}^{K} |\beta_k| \right). \quad (13)$$

There is also a regularization technique combining ridge regression and the LASSO called the elastic net regression. It is characterized by a penalty of the form $\lambda \sum_{k=1}^{K} \left(\alpha \beta_k^2 + (1 - \alpha) |\beta_k| \right)$, that is, a linear combination of the ridge-type and the LASSO-type penalties. Elastic net was the choice of regularization methods in Bloznelis (2016c).

The tuning parameter λ in ridge regression and the LASSO and α in the elastic net are normally found by cross validation. Cross validation is a generalization of the rolling window approach beyond the time series setting. Cross validation amounts to splitting the original sample into training and validation subsamples multiple times, estimating the candidate models on the training data and evaluating their performance on the validation data. A number of different λ and α values are tried out, and the values producing the best validation results are selected.

A good introduction into regularization methods is given in Hastie et al. (2009, p. 61-94) who also provides performance comparisons of subset selection to regularization and regularization to derived input directions methods. The performance of the different methods vary with the nature of the data generating processes, so that none of the methods is universally better or worse than all the remaining methods. Both the subset selection and regularization were used in Bloznelis (2016c). The use of regularization was relatively successful as the model yielding the highest forecast accuracy was a regularized model. Meanwhile, derived input directions, e.g. principal component analysis, has not been used due to space limitations.

5.5.2 Computational complexity

An important factor to consider in the context of model and/or variable selection is computational complexity. Often it may be the deciding factor as certain methods are prohibitively computationally expensive or completely infeasible. I will compare the computational expenses of subset selection and regularization techniques. Full subset selection is extremely expensive with the number of models to be fit growing exponentially in the number of variables (2^K for K variables). Estimation of each individual model has the complexity of a regular OLS regression fit, that is, $O(K^2 n)$, where n is the sample size. Hence, the complexity of full subset selection is $O(2^K K^2 n)$. However, brute-force complexity may be improved upon in the case of full subset selection; there exist fast algorithms making full subset selection of up to around 40 variables feasible on a modern personal computer; see e.g. Lumley and Miller (2009). Stepwise selection is considerably less demanding with up to K models to be fit given K candidate variables; hence, its complexity is $O(K^3 n)$. Meanwhile, ridge or LASSO regularization for a given penalty parameter λ is relatively efficient; computational complexity of ridge regression is the same as that of OLS regression; complexity of LASSO is $O(K^3 + K^2 n)$ if done using the LARS algorithm, which is again the same as that of OLS regression when $K < n$ (Efron et al., 2004). Since the penalty parameter is typically found from the data using cross validation, the computational complexity increases almost linearly in the number of sample splits, or folds (say, F), in cross validation; in Bloznelis (2016c) the number was about one hundred. For a λ grid of size L, running a ridge regression or a LASSO regression with tuning λ has the complexity of $O(LFK^2 n)$. The elastic net has an extra tuning parameter α that is also found using cross validation by examining the performance of a model with α being one of a number of grid points between zero and unity. Given a grid density of 0.1, the elastic net will be eleven times more computationally expensive than the LASSO or the ridge estimation. More generally, an elastic regression with α grid size A has the complexity of $O(ALFK^2 n)$. Table 2 reports the computational complexities of brute-force full subset selection, stepwise selection, ridge, LASSO and elastic net with respect to K and n. It shows that all the models considered have the same complexity with respect to sample size. However, full subset selection has exponential complexity in the number of parameters while the other methods have only polynomial complexity, which is superior.

Estimation speed was less of a concern for models estimated by OLS such as the VECM with a given cointegration vector. Meanwhile, it was considerably more important for models estimated by the maximum likelihood method such as ARIMA, regression with ARMA errors, and DCC-GARCH, because maximum likelihood estimation is typically more computationally intensive than OLS.

5.5.3 The case of non-regression-based models

The three categories of variable selection techniques discussed above are suited for models based on linear regression and are estimable using OLS. Meanwhile, time series models that do not take the form of a regular regression, have latent variables and are estimated using the maximum likelihood method (e.g. ARIMA, regression with ARMA errors, and DCC-GARCH) do not lend themselves directly to the aforementioned regularization techniques or to the methods using derived input directions; however, subset selection techniques are still applicable. Full subset selection was utilized with these models in Bloznelis (2016b) and Bloznelis (2016c). This technique avoids missing the most relevant subsets at the expense of high computational costs and proneness to picking up noise in place of signal. Bloznelis (2016a) uses an algorithm by Hyndman and Khandakar (2008) for stepwise ARIMA model selection which delivers a good balance between exhaustiveness and robustness. For completeness it should be noted that there exists a counterpart of regularization based on penalized likelihood estimation applicable to models estimated using the maximum

Table 2 Approximate computational complexity of subset selection and regularization techniques

Approximate computational complexity	Full subset selection (brute force)	Stepwise subset selection	Ridge regression	LASSO regression	Elastic net regression
With respect to sample size, n	$O(n)$	$O(n)$	$O(n)$	$O(n)$	$O(n)$
With respect to number of variables, K	$O(2^K)$	$O(K^3)$	$O(K^2)$	$O(K^3)$	$O(K^3)$

likelihood method; see e.g. Fan and Li (2001); also, an ARMA model could be approximated with a pure AR model which is estimable by OLS, and regularization could be used then.

In sum, model and variable selection plays an important role in this thesis, as it does in statistical modelling in general. This section has served as a brief introduction into an otherwise wide and deep field of statistical and machine learning approach to predictive model building, the results in which are widely applicable within and beyond economics.

5.6 Software implementation: R

The software implementation of the models in Bloznelis (2016b), Bloznelis (2016c) and Bloznelis (2016a) was done entirely in R (R Core Team, 2015). The choice of this particular statistical software package was motivated by a number of factors. (I will only list the factors that I encountered first-hand; I do not intend to give a comprehensive overview.) First, R has a broad coverage of the relevant models and estimation techniques. Second, it is open source which allows seeing how exactly the different routines are implemented and thus lowers the risk of misinterpretation and misuse of the built-in functionality. Third, a large and active user community effectively and efficiently substitutes for an official customer service; solutions to frequently encountered problems often can be found online at the R Mailing List (R Core Team), Stack Overflow (Stack Exchange Inc.) and Cross Validated (Stack Exchange Inc.). Fourth, continuous although irregular updates and bug fixes ensure there being no need to wait for a scheduled release of a new version; responsive authors of the different packages (libraries) are often very quick to fix any reported problems. Fifth, R is free and easy to get wherever needed, which effectively makes it portable. That helps avoid the problem of only having one licensed copy at one computer which might not always be accessible. It also makes the research easily reproducible without the need of purchasing commercial software. Since the papers in this thesis are all of applied nature, making the source code accessible to the salmon market participants could be an effective way of realizing the underlying commercial potential.

There are also a few disadvantages to R. First, it does not use multiple cores of the computer processor by default; hence, only a fraction of the theoretically available computing power is used at any time. This can be circumvented by using designated libraries for parallelizing the computations; however, the task may be burdensome and not completely straightforward for inexperienced users. Second, R is not very efficient in memory management and thus may require unnecessarily large amount of random access memory (RAM) and/or increased effort in writing memory-efficient programmes. These disadvantages have to be taken into consideration if computing power and/or memory availability is of concern.

6 Summary of the papers

Paper 1: Salmon price volatility: a multivariate weight-class-specific approach

Salmon price is highly volatile both in absolute terms and as compared to other commodities, and the high volatility is negatively perceived by the market participants. In this paper I analyze the development of the spot price volatility for different weight classes over time. The data used are spot prices for seven different weight classes as well as the average spot price for 1995-2013 and the EUR/NOK currency exchange rate. Univariate ARMA-GARCH models with exogenous regressors and dynamic conditional correlation (DCC) model are used to describe the behaviour of the conditional means and variances of the series.

The main findings are the following. An increase in the unconditional variance around 2006 is detected in all series, which may have occurred due to a change in industry regulations (introduction of the maximum allowable biomass restriction) or the introduction of the futures market for salmon. Volatility over 1996-2005 is considerably lower than in the 2007-2013 period; the persistence of volatility is low throughout the sample. The price dynamics of the conditional means and variances are similar across the neighbouring weight classes (fish of similar size). Conditional correlations across the weight classes are high, more so in the latter period and especially between fish of similar size. The EUR/NOK rate is found to be relevant for the conditional mean in 2007-2013; a 1% increase in the exchange rate corresponds to a 0.6% increase in the spot price. However, it is not found to significantly influence the conditional variance of price returns on any of the weight classes.

Strong similarities in the price dynamics and high conditional correlations between the neighbouring weight classes suggest that the three most popular weight classes (3-6 kg), accounting for around 75% of the total production volume, can be treated as one when it comes to price behaviour. The high conditional correlations, especially since 2007, indicate limited opportunities to hedge by substituting between the weight classes. This is conducive to the functioning of the futures market for salmon in two respects: it decreases the number of alternatives as to how hedging can be done; and it allows for greater hedging effectiveness if a mix of weight classes is hedged at once.

Paper 2: Short-run salmon price forecasting

Forecasting the spot price of salmon is a key activity of the salmon market participants. Whether to maximize revenue from selling the fish or minimize costs from purchasing it, forecasts of salmon price are used in business management by salmon farmers, processors, exporters and retailers. However, accurate price forecasts may be difficult to obtain. The aim of this paper is to forecast the average spot price of salmon one to five weeks ahead using different methods and evaluate the forecast accuracy and optimality. The variables considered relevant for salmon price forecasting are lagged spot price of salmon, production volume, salmon futures prices, share prices of salmon farming companies listed on the Oslo Stock Exchange, and the EUR/NOK currency exchange rate. The sample covers 2007-2014 (404 weekly observations) where the 2013-2014 period is used for an out-of-sample assessment of forecasting performance. Nine individual models with variations in their estimation techniques are explored and two forecast combination techniques are employed. The individual models are: (1) random walk; (2) seasonal random walk; (3) ARIMA; (4) ARFIMA; (5) ARARMA; (6a) VAR with lag orders determined via full subset selection; (6b) VAR estimated using elastic net regularization; (7a) VECM with lag orders determined via full subset selection; (7b) VECM estimated using elastic net regularization; (8) a model based on the k-nearest neighbours method; (9) artificial neural networks. The forecast combination techniques are arithmetic mean and trimmed arithmetic mean.

Model prediction for 2013-2014 confirm the previous experience that forecasting salmon price is extremely challenging. Although there is a predictable seasonal component, it barely helps reduce the forecast error as compared to a naïve forecast. The best forecasting model is (7a), i.e. VECM estimated using elastic net regularization. It yields the correct direction of price change in 50% - 64% of instances and Theil's U statistics in the range of 1,00 - 0,90, depending on the forecast horizon. The second best prediction method is (2) with its naïve seasonal forecast. The two methods' forecasts pass about half of the forecast optimality tests. Importantly, the economic value of the forecasts is limited in practice. The limited predictability beyond seasonality does not offer compelling evidence against weak form market efficiency of the salmon spot market.

Paper 3: Hedging salmon price risk

High volatility and limited predictability of the salmon spot price call for measures of risk management in the salmon spot market. This paper aims at designing a feasible strategy for hedging the future spot price of salmon and assessing its effectiveness. The spot price of salmon is to be hedged using salmon futures, live cattle futures, soybean meal and oil futures, and the share price of Marine Harvest on the Oslo Stock Exchange. The period of 2007-2015 (434 weekly observations) is used for in-sample modelling and out-of-sample performance assessment of 4-, 8- and 13-week hedges.

I propose a new measure of hedging effectiveness to match the objective of uncertainty minimization under partly predictable prices. The new measure is defined as relative reduction in the mean squared forecast error due to the hedging portfolio relative to the unhedged position. In the application, none of the financial instruments or their combinations appears to track the movement in the spot price of salmon closely enough to be practically relevant, except for the salmon futures. Hedging effectiveness of 0.12 through 0.58 is achieved for the average spot price and 0.07 through 0.60 for the specific weight classes, depending on the hedging horizon and the model for the conditional mean and variance of the hedging portfolio. As expected, the effectiveness increases with the hedging horizon. The best way of utilizing the futures contract is keeping it until maturity, which yields higher hedging effectiveness and lower transaction costs than closing the futures and the spot positions simultaneously. However, limited liquidity of the salmon futures market poses a problem for practical applications of the hedging strategy. Hence, there is still a need to look for new efficient and feasible ways to cope with the high uncertainty over the future spot price of salmon. For example, proliferation of forward contracting could serve the purpose. Also, the gradual vertical integration in the industry will likely make the high uncertainty less of a problem for an increasing number of market participants over time.

7 Contributions and limitations

The contributions of the research papers can be classified into methodological and empirical ones. Due to the empirical nature of the thesis, most of the contributions are empirical; also, the methodological contributions are closely tied to practical applications.

7.1 Methodological contributions

The first methodological contribution is stressing the peculiarity of seasonal adjustment for weekly data and use of appropriate seasonal adjustment procedures. Weekly data is characterized by its non-integer frequency as there are a fractional number of weeks in a year. The widespread practice of assuming the number of weeks to equal 52 is simplistic and generally not suitable for modelling time series spanning more than just two or three years. A few different approaches to weekly seasonality modelling have been proposed by Harvey et al. (1997), Cleveland and Scott (2007), Hyndman (2010b) and Hyndman (2014), and others. The approach in the latter two sources is simple and easy to implement, yet it sensibly addresses the major problem of the fractional frequency. I have chosen to follow this approach for its flexibility and simplicity and used it in all of the three papers.

Second, Bloznelis (2016a) proposes a new measure of hedging effectiveness, namely the relative reduction in mean squared forecast error, defined as

$$RRMSFE = \frac{E_t\left(\left(s_{t+h} - E_t(s_{t+h})\right)^2\right) - E_t\left(\left((s_{t+h} - \beta f_{t+h}) - E_t(s_{t+h} - \beta f_{t+h})\right)^2\right)}{E_t\left(\left(s_{t+h} - E_t(s_{t+h})\right)^2\right)} \tag{14}$$

where s_t is the asset price at time t, f_t is the hedging instrument's price, β is the hedge ratio and $E_t(\cdot)$ denotes conditional expectation given the information available at time t. The measure is specifically suited for assets with predictable prices where either (1) material gains or losses from the predicted price changes cannot be realized due to perishability of the asset or (2) the hedger is insensitive to the level of returns, or both. It is price predictability that suggests the hedger's objective function differs from the one normally used in financial applications which are characterized by unpredictability of the conditional mean. The new measure of hedging effectiveness follows naturally from the objective function of uncertainty minimization under square loss.

Third, Bloznelis (2016a) examines a hedging strategy that is often overlooked in empirical studies, namely, holding a futures contract through maturity – rather than liquidating the spot and the financial positions simultaneously. The strategy appears to be the most effective among its competitors for hedging the uncertainty in the future spot price of salmon. Moreover, it is also relatively cost effective and easy to implement for practitioners; the avoidance of closing the contract before maturity saves half the commission costs and half the trading efforts as compared to the simultaneous liquidation strategy.

7.2 Empirical contributions

In the following I present this thesis' empirical contributions relevant for the salmon market participants and/or academic researchers.

First, Bloznelis (2016b) finds increased conditional correlations between prices of different weight classes of salmon in 2007-2013, suggesting that salmon has become more homogenous across the weight classes. This shrinks the potential for diversification using different weight classes (to which there are also biological limits). In turn, this is conducive to the functioning of the futures market for salmon in the following two respects. First, the unmet demand for hedging may shift away from weight-class diversification and plausibly towards hedging using the salmon futures. Second, salmon futures has become a better hedging instrument for the different weight classes as the price dynamics across the weight classes turned more homogenous and closer to the underlying of the futures contracts.

Second, Bloznelis (2016c) provides an example where forecast combinations do not outperform individual forecasts. This is in contrast to the typical findings and may merit a study on the conditions that yield this phenomenon.

Third, Bloznelis (2016c) finds the spot price to be essentially unpredictable beyond seasonality. This provides an indication for where the attention could be focused in the future studies of salmon price forecasting. An accurate seasonal forecast may be the key to an accurate price forecast. Unpredictability beyond seasonality also suggests weak form efficiency in the spot market for salmon; the finding is relevant for the market participants' pricing strategies and formation of the price expectations.

Fourth, Bloznelis (2016a) demonstrates how optimal hedge ratios could be calculated given the newly proposed objective function of a hedger. The paper discusses the values and variability of the optimal hedge ratios in the recent years for hedging the spot price of salmon with salmon futures. This might provide a useful starting point for a commercial hedger.

Fifth, Bloznelis (2016a) quantifies the potential uncertainty reduction due to hedging based on the recent data. Relative reduction in the mean squared forecast error is documented for different weight classes and different hedging strategies using salmon futures contracts. Cross hedging with live cattle, soybean meal and oil futures and the share price of Marine Harvest is shown to be inauspicious.

Sixth, the three research papers together provide a broad overview of options of uncertainty management available for the Norwegian farmed Atlantic salmon industry. This should enhance the practitioners' awareness and understanding of the different strategies and their expected effects.

7.3 Limitations

Due to the nature of the research problems, the thesis does not focus on causal explanations. While the choice of questions to ask and problems to solve can hardly be classified as a serious limitation of a study, it could have nevertheless been interesting to examine causality of the observed phenomena. Hence, a more structural approach to spot price forecasting and hedging may merit academic studies in the future.

A necessary minimum of assumptions is taken on in the three papers. Most of the models used are linear, which is a common simplification in the economics literature. Bloznelis (2016c) and Bloznelis (2016a) use the rolling-window approach which should lend robustness to the findings and essentially provides a robustness check in itself; for example, Bloznelis (2016a) shows the variation in the optimal hedge ratios over time. Model selection is fully automated in all the three papers, following explicit and concrete principles; hence, the influence of human judgement and the inherent tendency of selective reporting (Ioannidis, 2005) is limited. Modelling procedures being automated raises hope that the modelling approach could continue yielding similar results as those presented in the research papers, absent structural changes in the data generating processes. There remains a question whether findings stemming from a predictive modelling approach and reduced-form time series models are inherently more or less robust than those from explanatory approach and structural time series models. The key difference is absence or presence of theory-implied assumptions on the models. As long as the assumptions closely reflect the reality, they may be expected to robustify the modelling process and the results. However, once inappropriate assumptions are taken upon, the resultant inflexibility may prevent finding more appropriate models.

There are limitations on the applicability of the findings in practice. The focus on the short-term phenomena makes the results relatively sensitive to ageing. However, this is more or less true in all empirical studies. Crucially, the approaches taken in the papers are morally up to date and will lend themselves directly to repeated applications on any new data to come. On the other hand, it is noteworthy that advanced knowledge of econometrics and statistics supplemented by sufficient experience with statistical software will be needed to replicate the analysis. The latter may be perceived as a natural prerequisite for success in today’s data-rich and technologically capable world.

References

ANDERSEN, T. B., ROLL, K. H. & TVETERÅS, S. 2008. The price responsiveness of salmon supply in the short and long run. *Marine Resource Economics,* 23**,** 425-437.

ASCHE, F. 1996. A system approach to the demand for salmon in the European Union. *Applied Economics,* 28**,** 97-101.

ASCHE, F. & BJORNDAL, T. 2011. *The economics of salmon aquaculture,* Chichester, John Wiley & Sons.

ASCHE, F., BJØRNDAL, T. & SALVANES, K. G. 1998. The demand for salmon in the European Union: the importance of product form and origin. *Canadian Journal of Agricultural Economics/Revue canadienne d'agroeconomie,* 46**,** 69-81.

ASCHE, F., BREMNES, H. & WESSELLS, C. R. 1999. Product aggregation, market integration, and relationships between prices: an application to world salmon markets. *American Journal of Agricultural Economics,* 81**,** 568-581.

ASCHE, F., GUTTORMSEN, A. G., SEBULONSEN, T. & SISSENER, E. H. 2005. Competition between farmed and wild salmon: The Japanese salmon market. *Agricultural Economics,* 33**,** 333-340.

ASHEIM, L. J., DAHL, R. E., KUMBHAKAR, S. C., ØGLEND, A. & TVETERÅS, R. 2011. Are prices or biology driving the short-term supply of farmed salmon? *Marine Resource Economics,* 26**,** 343-357.

BAUWENS, L., LAURENT, S. & ROMBOUTS, J. V. 2006. Multivariate GARCH models: a survey. *Journal of Applied Econometrics,* 21**,** 79-109.

BERGFJORD, O. J. 2007. Is there a future for salmon futures? An analysis of the prospects of a potential futures market for salmon. *Aquaculture Economics & Management,* 11**,** 113-132.

BIRD, P. J. 1986. Econometric estimation of world salmon demand. *Marine Resource Economics,* 3**,** 169-182.

BJØRNDAL, T., SALVANES, K. G. & GORDON, D. V. 1994. Elasticity estimates of farmed salmon demand in Spain and Italy. *Empirical Economics,* 19**,** 419-428.

BLOOMBERG. 2010. *Target Eliminates Farmed Salmon From All Target Stores* [Online]. Bloomberg. Available: http://www.bloomberg.com/apps/news?pid=newsarchive&sid=ahGdJIivWrzI [Accessed 2015-10-12].

BLOZNELIS, D. 2016a. Hedging salmon price risk. *Preprint.*

BLOZNELIS, D. 2016b. Salmon price volatility: a weight-class-specific multivariate approach. *forthcoming in Aquaculture Economics & Management*.

BLOZNELIS, D. 2016c. Short-term salmon price forecasting. *Preprint.*

BOEHLKE, J. 2015. *The Health Benefits of Salmon* [Online]. Livestrong.com. Available: http://www.livestrong.com/article/117686-health-benefits-salmon/ [Accessed 2015-10-12].

BOLLERSLEV, T. 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics,* 31**,** 307-327.

BOLLERSLEV, T. Glossary to ARCH (GARCH). *In:* BOLLERSLEV, T., RUSSELL, J. R. & WATSON, M. W., eds. Volatility and time series econometrics: essays in honour of Robert F. Engle, 2009. Oxford University Press.

BOLLERSLEV, T., ENGLE, R. F. & WOOLDRIDGE, J. M. 1988. A capital asset pricing model with time-varying covariances. *The Journal of Political Economy,* 96**,** 116-131.

BOX, G. E. & JENKINS, G. M. 1970. *Time series analysis: forecasting and control,* San Francisco, Holden Day.

BOX, G. E., JENKINS, G. M. & REINSEL, G. C. 2011. *Time series analysis: forecasting and control*, John Wiley & Sons.

BREIMAN, L. 2001. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical Science,* 16**,** 199-231.

CAPORIN, M. & MCALEER, M. 2013. Ten things you should know about the dynamic conditional correlation representation. *Econometrics,* 1**,** 115-126.

CLEVELAND, W. P. & SCOTT, S. 2007. Seasonal adjustment of weekly time series with application to unemployment insurance claims and steel production. *Journal of Official Statistics,* 23**,** 209-221.

COASTAL ALLIANCE FOR AQUACULTURE REFORM. *Salmon Farming Problems* [Online]. Available: http://www.farmedanddangerous.org/salmon-farming-problems/ [Accessed 2015-10-12].

CONT, R. 2001. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance,* 1**,** 223-236.

EFRON, B., HASTIE, T., JOHNSTONE, I. & TIBSHIRANI, R. 2004. Least angle regression. *The Annals of Statistics,* 32**,** 407-499.

ENGLE, R. F. 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica: Journal of the Econometric Society,* 50**,** 987-1007.

ENGLE, R. F. 2002. Dynamic conditional correlation: a simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics,* 20**,** 339-350.

ENGLE, R. F. & SHEPPARD, K. 2001. Theoretical and empirical properties of dynamic conditional correlation multivariate GARCH. *NBER Working Papers.* Cambridge: National Bureau of Economic Research.

EWALD, C.-O. 2013. Derivatives on nonstorable renewable resources: fish futures and options, not so fishy after all. *Natural Resource Modeling,* 26**,** 215-236.

EWALD, C.-O., NAWAR, R., OUYANG, R. & SIU, T.-K. 2014. Salmon futures and the Fish Pool market: the relevance of the Schwartz 97 two-factor model for fish farming. *Available at SSRN 2136965*.

EWALD, C.-O. & SALEHI, P. 2015. Salmon futures and the Fish Pool market in the context of the CAPM and the Fama & French three-factor model. *Available at SSRN*.

FAN, J. & LI, R. 2001. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association,* 96**,** 1348-1360.

FARMED SALMON BOYCOTT. 2015. *Stopping farmed salmon at the cash register* [Online]. Available: http://farmed-salmon-boycott.com/ [Accessed 2015-10-12].

FISH POOL. 2014. *NOS and NASDAQ price* [Online]. Fish Pool. Available: http://fishpool.eu/wp-content/uploads/2014/08/FAQ-Nasdaq.pdf [Accessed 2015-03-30].

FISH POOL. 2015a. *Appendix 1 to Fish Pool Rulebook* [Online]. Available: http://fishpool.eu/wp-content/uploads/2015/03/Fish-Pool-Rulebook-App-1.pdf [Accessed 2015-10-13].

FISH POOL 2015b. Arsrapport 2014. Oslo/Bergen/Florø.

FISH POOL. 2015c. *Fish Pool Index* [Online]. Available: http://fishpool.eu/price-information/spot-prices/fish-pool-index/ [Accessed 2015-10-12].

GRINDHEIM, J. 2009. *Laksepris på smertegrensa* [Online]. Available: http://www.intrafish.no/norsk/nyheter/article1296227.ece [Accessed 2015-10-12].

GRINDHEIM, J. 2013. *Farmed salmon prices see recovery* [Online]. Available: http://www.intrafish.com/news/article1378140.ece [Accessed 2015-10-12].

GU, G. & ANDERSON, J. L. 1995. Deseasonalized state-space time series forecasting with application to the US salmon market. *Marine Resource Economics,* 10**,** 171-185.

GUTTORMSEN, A. G. 1999. Forecasting weekly salmon prices: risk management in fish farming. *Aquaculture Economics & Management,* 3**,** 159-166.

HAMILTON, J. D. 1994. *Time series analysis,* Princeton, Princeton University Press.

HANSEN, P. R. & LUNDE, A. 2005. A forecast comparison of volatility models: does anything beat a GARCH (1, 1)? *Journal of Applied Econometrics,* 20**,** 873-889.

HARVEY, A., KOOPMAN, S. J. & RIANI, M. 1997. The modeling and seasonal adjustment of weekly observations. *Journal of Business & Economic Statistics,* 15**,** 354-368.

HASTIE, T., TIBSHIRANI, R. & FRIEDMAN, J. 2009. *The elements of statistical learning,* Berlin, Springer.

HYNDMAN, R. J. 2010-10-04 2010a. The ARIMAX model muddle. *Hyndsight* [Online]. Available from: http://robjhyndman.com/hyndsight/arimax/ [Accessed 2015-10-04].

HYNDMAN, R. J. 2010b. Forecasting with long seasonal periods. *Hyndsight* [Online]. Available from: http://robjhyndman.com/hyndsight/longseasonality/ [Accessed 2015-10-13].

HYNDMAN, R. J. 2014. Forecasting weekly data. *Hyndsight* [Online]. Available from: http://robjhyndman.com/hyndsight/forecasting-weekly-data/ [Accessed 2014-03-05].

HYNDMAN, R. J. 2015. The bias-variance decomposition. *Hyndsight* [Online]. Available from: http://robjhyndman.com/hyndsight/bias-variance/ [Accessed 2015-08-03].

HYNDMAN, R. J. & ATHANASOPOULOS, G. 2014. Forecasting: principles and practice. OTexts.

HYNDMAN, R. J. & KHANDAKAR, Y. 2008. Automatic time series for forecasting: the forecast package for R. *Journal of Statistical Software,* 27**,** 1-22.

HOERL, A. E. 1962. Application of ridge analysis to regression problems. *Chemical Engineering Progress,* 58**,** 54-59.

HOERL, A. E. & KENNARD, R. W. 1970. Ridge regression: biased estimation for nonorthogonal problems. *Technometrics,* 12**,** 55-67.

IOANNIDIS, J. P. 2005. Why most published research findings are false. *Chance,* 18**,** 40-47.

IPATENCO, S. *Why salmon is one of the best fishes you can eat* [Online]. SFGate. Available: http://healthyeating.sfgate.com/salmon-one-fishes-can-eat-5512.html#page1 [Accessed 2015-10-12].

JENSEN, B.-A. 2011. *Analysts: surprise and uncertainty over salmon prices* [Online]. Available: http://www.intrafish.com/news/article1257157.ece [Accessed 2015-10-12].

JOHANSEN, S. 1995. *Likelihood-based inference in cointegrated vector autoregressive models,* Oxford, Oxford University Press.

LEWIN, J. 2014. *The health benefits of... salmon* [Online]. Available: http://www.bbcgoodfood.com/howto/guide/ingredient-focus-salmon [Accessed 2015-10-12].

LIN, B. H., HERRMANN, M., LIN, T. Y. & MITTELHAMMER, R. C. 1989. Forecasting the price of farmed Atlantic salmon: an integrated econometric and time series approach. *Agribusiness,* 5**,** 477-488.

LIU, Y., OLAUSSEN, J. O. & SKONHOFT, A. 2011. Wild and farmed salmon in Norway—a review. *Marine Policy,* 35**,** 413-418.

LUMLEY, T. & MILLER, A. 2009. leaps: regression subset selection. 2.9 ed. Vienna, Austria: R Project for Statistical Computing.

LÜTKEPOHL, H. 2007. *New introduction to multiple time series analysis,* Berlin, Springer.

MARINE HARVEST 2015. Salmon farming industry handbook 2015. Bergen: Marine Harvest.

NILSEN, G. B. 2011. *Salmon expert cannot explain price drop* [Online]. Available: http://www.intrafish.com/news/article1345231.ece [Accessed 2015-10-12].

NORWEGIAN SEAFOOD FEDERATION & NORWEGIAN SEAFOOD COUNCIL. 2011. *Aquaculture in Norway* [Online]. Available: http://sjomatnorge.no/wp-content/uploads/importedfiles/Aquaculture%2520in%2520Norway%25202011.pdf [Accessed 2015-10-12].

OGLEND, A. 2013. Recent trends in salmon price volatility. *Aquaculture Economics & Management,* 17**,** 281-299.

OGLEND, A. & SIKVELAND, M. 2008. The behaviour of salmon price volatility. *Marine Resource Economics,* 23**,** 507-526.

OWEN, J. 2008. *Farmed salmon decimating wild salmon worldwide* [Online]. National Geographic News. Available: http://news.nationalgeographic.com/news/2008/02/080212-salmon-lice.html [Accessed 2015-10-12].

R CORE TEAM. *R mailing list* [Online]. Vienna, Austria: R Project for Statistical Computing. Available: https://www.r-project.org/mail.html [Accessed 2016-01-27].

R CORE TEAM 2015. R: a language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing.

SHMUELI, G. 2010. To explain or to predict? *Statistical Science,* 25**,** 289-310.

SILVENNOINEN, A. & TERÄSVIRTA, T. 2009. Multivariate GARCH models. *In:* ANDERSEN, T. G., DAVIS, R. A., KREISS, J.-P. & MIKOSCH, T. V. (eds.) *Handbook of financial time series.* Berlin: Springer.

SOLIBAKKE, P. B. 2012. Scientific stochastic volatility models for the salmon forward market: forecasting (un-) conditional moments. *Aquaculture Economics & Management,* 16**,** 222-249.

STACK EXCHANGE INC. *Cross Validated* [Online]. Delaware: Stack Exchange Inc. Available: http://stats.stackexchange.com/ [Accessed 2016-01-27].

STACK EXCHANGE INC. *Stack Overflow* [Online]. Stack Exchange Inc. Available: http://stackoverflow.com/ [Accessed 2016-01-27].

TERÄSVIRTA, T. 2009. An introduction to univariate GARCH models. *In:* ANDERSEN, T. G., DAVIS, R. A., KREISS, J.-P. & MIKOSCH, T. V. (eds.) *Handbook of financial time series.* Berlin: Springer.

TIBSHIRANI, R. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological),* 58**,** 267-288.

VUKINA, T. & ANDERSON, J. L. 1994. Price forecasting with state-space models of nonstationary time series: case of the Japanese salmon market. *Computers & Mathematics with Applications,* 27**,** 45-62.

WHITTLE, P. 1951. *Hypothesis Testing in Time Series Analysis,* New York, Hafner.

XIE, J., KINNUCAN, H. W. & MYRLAND, Ø. 2009. Demand elasticities for farmed salmon in world trade. *European Review of Agricultural Economics,* 36**,** 425-445.

XIE, J. & MYRLAND, Ø. 2011. Consistent aggregation in fish demand: a study of French salmon demand. *Marine Resource Economics,* 26**,** 267-280.

Paper 1

Salmon price volatility:

a weight-class-specific multivariate approach

Daumantas Bloznelis[1]

School of Economics and Business

Norwegian University of Life Sciences

Abstract

Salmon spot price has been highly volatile and hard to predict since mid-2000s, obscuring the industry players' planning decisions. ARMA-GARCH and dynamic conditional correlation (DCC) models were employed on weekly data for 1995-2013 to examine the behavior of weight-class-specific prices which are directly relevant for salmon production and risk management. Two periods of different volatility regimes were identified, before and after 2006 when the salmon market was undergoing fundamental changes. Both volatility and conditional correlations increased from 1996-2005 to 2007-2013, and return dynamics became more homogenous across weight classes. This development is conducive to the functioning of the salmon futures and options exchange.

Keywords: salmon market, price volatility, GARCH, dynamic conditional correlation

[1] Mailing address: School of Economics and Business, Norwegian University of Life Sciences, P.O. Box 5003, NO-1432 Ås, Norway. Phone: (47) 64 96 64 62. E-mail: daumantas.bloznelis@nmbu.no.

1 Introduction

The market for the Norwegian farmed Atlantic salmon is characterized by high price volatility. The standard deviation of logarithmic returns (returns) on weekly average salmon spot price more than doubled from 3.0% in 1996-2005 to 7.3% in 2007-2013. Compared with feeder cattle, wheat and other agricultural commodities, salmon price volatility was relatively low in the former period but exceptionally high in the latter one.

High price volatility is generally considered negative, especially for producers who are averse to risk; see e.g. Jensen (2013). This is not surprising since a large share of salmon transactions take place in the spot market without price hedging. Salmon processors are struggling to maintain profitability consistently given the highly volatile price of raw salmon, which is their production input, while selling processed salmon at relatively “sticky” prices (Bergfjord, 2007). Processors are thus absorbing a large portion of the price uncertainty. Understanding the development of price and its volatility over time could help processors in their production and price risk management decisions. Salmon farmers are facing price uncertainty when selling their production. They would benefit from slaughter-timing strategies targeted at maximizing the price and minimizing uncertainty. Meanwhile, suboptimal slaughtering schedule could be especially costly in an environment of high price volatility. Thus both processors and farmers could profit from better understanding the dynamics of the price volatility.

One strategy for managing price uncertainty is hedging. The opening of “Fish Pool” futures and options exchange in 2006 made hedging via futures contracts possible. However, only up to 10% of the total Norwegian production volume is being hedged at “Fish Pool”. Hedging can also be carried out by using bilateral forward contracts. Although the volume of such bilateral hedging is unknown to me, Larsen and Asche (2011) claim that the use of bilateral contracts has increased substantially in the recent decades.

Considering the origins and development of salmon price volatility, there are a number of relevant factors to study. First, variation in price is likely associated with variation in production volume. After a sudden change in slaughter volume that is not offset by a corresponding change in demand, we may expect a change in price so that supply and demand would be equalized. For example, slaughter volume may increase due to forced harvesting following an outbreak of a disease. It may decrease due to bad weather conditions such as autumn and winter storms that occasionally interrupt the normal functioning of salmon farms on the west coast of Norway. Thus we may expect price to be more volatile when volume is more volatile, and vice versa. Oglend and Sikveland (2008) find this relationship to be statistically significant but economically small for the average spot price of salmon during the period of 1995-2007. Yet it is not clear what kind of volume data they use. While the total spot market volume data are not being collected, the total export volume data may be obtained from Statistics Norway (SSB). However, the latter need not be a good proxy for the former since the share of the spot market volume in the total export volume may vary considerably over time.

Second, volatility in the currency exchange rates may indirectly affect volatility in the salmon price. Norway exports almost all of its salmon production and the vast majority of export transactions are invoiced in foreign currencies (Straume, 2014).If there is price transmission from the export price to the spot price, exchange rates may be important drivers of the spot price measured in NOK. Accordingly, high volatility in the relevant exchange rates (e.g. EUR/NOK) could lead to high volatility in the salmon price. Oglend (2013) modelled this relationship, but found no significant effect of returns on the EUR/NOK rate on volatility of returns on the average salmon price.

Prices of substitutes for salmon, such as beef, pork and chicken; other demand factors; cost of production inputs; terms of trade; and other may also determine salmon price volatility. See also Oglend and Sikveland (2008) and Oglend (2013). The recent studies by Brækkan and Thyholdt (2014) and Brækkan (2014) find global salmon demand to be rather variable over 2002-2011, which may also provide part of the explanation for the high salmon price volatility.

Previous academic studies of salmon price volatility are few, considering the size and importance of the salmon farming industry. An early study by Oglend and Sikveland (2008) analyses weekly returns on the average spot price of salmon over 1995-2007. The authors apply a univariate autoregressive (AR) conditional mean model and a generalized autoregressive conditional heteroskedasticity (GARCH) model for the conditional variance. They find presence of volatility clustering, positive correlation between price level and volatility level, and negative correlation between change in production volume and price returns. They also note that the volatility properties of salmon price are largely in line with those of other agricultural and fish commodities. In an up-to-date paper, Oglend (2013) analyses the recent trends in salmon price volatility. He lists important industry events to explain the increase in volatility observed around 2006 and builds a univariate AR-GARCH model to explain the conditional mean and the conditional variance of returns on the average spot price. Links to food price indices are shown to be significant in the volatility model, where increasing values of the food price indices coincide with increasing salmon price volatility. Solibakke (2012) analyses returns on daily front month futures at “Fish Pool” for 2006-2010. With a stochastic volatility model, he estimates the annualized volatility of returns at 16.25%. Positive asymmetry (leverage effect) is found such that volatility associated with positive price shocks is larger than that associated with negative price shocks. Note that Oglend and Sikveland (2008) and Oglend (2013) focus on the spot price of salmon while Solibakke (2012) explores the futures price, which may not be directly comparable. A general descriptive analysis of fish price volatility can be found in Dahl and Oglend (2014).

We may broaden the understanding of volatility stemming from these articles by using more detailed data and answering more comprehensive questions. In practice, both salmon farmers and buyers distinguish between different fish sizes, the differences arising from production technology and schedule, processing technology and consumers' preferences. Annual cycles in seawater temperature result in fish gaining weight at different paces over the year. This has led salmon farmers to adopt seasonal strategies of when to put smolt (the young fish) into seawater and how intensively to feed them at a given time of the year. As a result, the weight-class distribution of slaughter-ready salmon is changing throughout the year. This translates to varying relative prices across the weight classes (Asche & Guttormsen, 2001; Marine Harvest, 2013). From the processor's perspective, Marine Harvest (2013) indicates that European processors mainly use 3-6 kg fish although niche markets exist for other sizes. Meanwhile, different groups of consumers may prefer larger or smaller salmon based on the intended use or simply as a matter of taste. Also, the degree of homogeneity across weight classes may be an important determinant of the well- or ill-functioning of the salmon futures and options exchange (Bergfjord, 2007). Higher correlation across weight classes makes the futures price more viable as a hedge for the less voluminous and thus "non-standard" weight classes. As a counterfactual illustration, Martinez-Garmendia and Anderson (1999) argue that severe price differences across different weight classes of shrimp may have caused the collapse in the white shrimp futures market. Caution is warranted as the latter example does not apply directly to our case; salmon futures are cash settled against a price index, while shrimp futures used to have actual delivery of the product. In sum, the properties of prices for specific weight classes are more directly relevant than those of the average price.

The focus of this study is on the price volatility for the specific weight classes of salmon. I set out to detect, describe and explain patterns in the development of volatility over time. I will identify differences between weight classes and explore pairwise and system-wide associations across them.

The next section provides an overview of the methods used in this study. This is followed by a discussion of the data. Econometric results are reported in the penultimate section, and a summary of findings concludes the paper.

2 Methods

Studies of volatility in the financial markets and subsequently in the commodity markets have revealed a set of stylized facts. Cont (2001) pins down the following asset return regularities: heavy-tailed conditional and unconditional distributions; gain versus loss asymmetry that produces skewed distributions; leverage effect as higher returns are associated with higher volatility; correlation between trading volume and volatility; absence of autocorrelations; slow decay of autocorrelation in absolute values of returns; volatility clustering; and other. Note that absence of autocorrelations need not be characteristic to commodities as opposed to pure financial assets.

2.1 (G)ARCH model

A number of stylized facts can be replicated by the (generalized) autoregressive conditional heteroskedasticity (G)ARCH type of models. Since the introduction of ARCH model by Engle (1982) and its extension to GARCH by Bollerslev (1986), various versions of the (G)ARCH model have become workhorses in volatility studies. The basic GARCH(m,s) model is defined as follows:

$$r_{t|t-1} = \sigma_{t|t-1}\varepsilon_t, \quad \varepsilon_t \sim i.i.d(0,1), \tag{1}$$

$$\sigma^2_{t|t-1} = \omega + \sum_{l=1}^{m} \alpha_l r^2_{t-l} + \sum_{k=1}^{s} \beta_k \sigma^2_{t-k|t-k-1}. \tag{2}$$

Here $r_{t|t-1}$ is the logarithmic price return at the time period *t* conditional on information available up to and including the time period *t-1*. Its time-varying conditional variance $\sigma^2_{t|t-1}$ evolves in a way that closely resembles an autoregressive moving average (ARMA) process, as evident from equation (2). Let us call the lagged returns “ARCH terms” and the lagged conditional variances “GARCH terms”.

The simple GARCH(1,1) specification

$$r_{t|t-1} = \sigma_{t|t-1}\varepsilon_t, \quad \varepsilon_t \sim i.i.d(0,1), \tag{3}$$

$$\sigma^2_{t|t-1} = \omega + \alpha_1 r^2_{t-1} + \beta_1 \sigma^2_{t-1|t-2} \tag{4}$$

is often found to be flexible enough to capture the stylized facts without the need to increase the orders m and s beyond unity. GARCH(1,1) describes a case where the return r_t has a basis level of variance ω and where the conditional variance at the time period t is a weighted sum of the basis level, the most recent squared return and the conditional variance of the previous period. Typical values of α_1 found for daily returns are close to zero, while β_1 is typically close to unity; for weekly or monthly data, α_1 may be larger while β_1 may be smaller. A relatively small value of α_1 would imply that the current return tends to have a small impact on the conditional variance of the next period's return. A relatively large value of β_1 signals that conditional variance typically remains close to its most recent value.

The sum of alphas and betas, $\lambda = \sum_{l=1}^{m} \alpha_l + \sum_{k=1}^{s} \beta_k$ in the case of GARCH(m,s) and $\lambda = \alpha_1 + \beta_1$ in the case of GARCH(1,1), is called persistence. The closer λ is to unity, the longer it takes for the conditional variance to revert to its long-term mean (the unconditional variance). Typically, λ is found to be almost equal to unity. If $\lambda = 1$, the conditional variance process is no longer mean reverting but integrated. Hence, the conditional variance can increase unboundedly without a tendency to return to its long-term mean. If $\lambda > 1$, conditional variance grows explosively. The last two cases are not welcome from the theoretical point of view, thus in practice parameter estimation is often accompanied by an inequality restriction $\lambda < 1$. I also follow this convention.

A useful way to measure the speed of mean reversion is to consider the time needed for the conditional variance to close half of the gap between its current value and its long-term mean. This time is called half-life, denoted K, and can be calculated as $K = \frac{ln(0.5)}{ln(\lambda)}$. For example, if $\lambda = 0.9$, $K = 6.6$, which means that half of the initial gap between the current conditional variance and its long-term mean is closed in less than seven periods. Obviously, for $\lambda = 1$, $K = \infty$, and the gap tends to remain for an infinitely long time.

Before estimating a (G)ARCH model, its relevance should be investigated. A common way to do that is to test for the so-called ARCH effects using the ARCH-LM test proposed by Engle (1982). Essentially, the test is conducted by estimating an ARCH(m) model and testing the model's overall significance. Since GARCH(m,s) can be expressed as ARCH(∞), the test not only captures the presence of ARCH patterns, but also possibly GARCH structures. The choice of lag length, m, is non-trivial and involves a trade-off between the test power against specific alternatives and its generality.

Extensions of (G)ARCH models may be applied in a multivariate context. For example, consider volatility of returns on a portfolio of assets. Each individual asset could be modelled using a univariate (G)ARCH model. However, a collection of univariate (G)ARCH models will typically not suffice to describe the volatility dynamics of the whole portfolio. The variance of portfolio returns is the sum of individual asset variances and twice the sum of covariances between them. Thus the whole covariance matrix of the asset portfolio is of interest unless the asset returns are uncorrelated.

A portfolio of salmon of different sizes (different weight classes) could be analyzed in a similar framework. Here, univariate (G)ARCH models for different weight classes' returns define their respective conditional variances, or the diagonal elements of the time-varying portfolio covariance matrix. Conditional covariances, or the off-diagonal elements, may be obtained using a dynamic conditional correlation (DCC) model by Engle (2002)[2].

2.2 DCC model

The DCC model considers conditional correlations (i.e. conditional covariances divided by conditional variances) between returns across different weight classes. Define standardized returns:

$$\eta_{i,t} = \frac{r_{i,t}}{\sigma_{i,t}} \tag{5}$$

[2] Other multivariate extensions of GARCH models are also available; for a review see Bauwens et al. (2006) or Silvennoinen and Teräsvirta (2009), among others. Elucidating critique of the DCC model is given in Caporin and McAleer (2013).

where i indexes the weight classes, t denotes the time period and $\sigma_{i,t}$ are obtained from univariate (G)ARCH models for each weight class separately. For a given pair of weight classes i and j, define conditional quasi-correlations of their respective returns at time t:

$$q_{ij,t|t-1} = (1-\gamma-\delta)\rho_{ij} + \gamma\eta_{i,t-1}\eta_{j,t-1} + \delta q_{ij,t-1|t-2} \tag{6}$$

The following terms govern the dynamics of the conditional quasi-correlation in equation (6): the regular (time-invariant) unconditional correlation ρ_{ij}; a product of lagged standardized returns defined in equation (5), or a moving-average term, $\eta_{i,t-1}\eta_{j,t-1}$; and an autoregressive component $q_{ij,t-1|t-2}$. Conditional correlation is just the scaled conditional quasi-correlation:

$$\rho_{ij,t|t-1} = \frac{q_{ij,t|t-1}}{\sqrt{q_{ii,t|t-1}}\sqrt{q_{jj,t|t-1}}} \tag{7}$$

The main difference between correlations and quasi-correlations is that correlations are always bounded between -1 and 1 while quasi-correlations may exceed these bounds.

We may distinguish two parts of equation (6): the static, or unconditional, part consisting of the first term; and the dynamic, or conditional, part that is time-varying and consists of the latter two terms. Weights $1-\gamma-\delta$ and $\gamma+\delta$ show the relative importance of the static and the dynamic part, respectively. By using conventional measures of joint parameter significance, it is possible to test whether the dynamic part should be included in the model against a static-only model. The DCC model provides a parsimonious way of describing the dynamics in conditional covariances and is easy to compute even in large systems of returns (Engle, 2002).

For simplicity of exposition, I have so far only considered the basic DCC(1,1) model. Generalization to a DCC model with higher lag orders is straightforward; more lags of the moving average and/or the autoregressive term could be added in equation (6). Furthermore, an extra term may be allowed in equation (6) to account for asymmetric reaction of conditional quasi-correlation to positive versus negative lagged returns (Cappiello et al., 2006):

$$q_{ij,t|t-1} = (1-\gamma-\delta)\rho_{ij} + \gamma\eta_{i,t-1}\eta_{j,t-1} + \theta neg\big(\eta_{i,t-1}\big)neg(\eta_{j,t-1}) + \delta q_{ij,t-1|t-2} \tag{8}$$

where the $neg(\cdot)$ operator sets any non-negative values to zero but leaves all negative values unchanged. Again, the model can easily be generalized to allow for greater lag orders.

Another source of flexibility in the dynamics of conditional correlations may be added by allowing the DCC model coefficients to differ across blocks of portfolio elements. Suppose the whole set of return series can be grouped into a few blocks. The assignment of portfolio elements to blocks would have economic interpretation. E.g., when the portfolio consists of series of financial returns, blocks could be formed by grouping the series sector-wise (financials, utilities, consumer staples etc.). Then the flexible DCC model by Billio et al. (2006) is obtained:

$$q_{ij,t|t-1} = (1-\gamma-\delta)\rho_{ij} + \gamma_{block(i)}\gamma_{block(j)}\eta_{i,t-1}\eta_{j,t-1} + \delta_{block(i)}\delta_{block(j)}q_{ij,t-1|t-2} \tag{9}$$

Here the subscript $block(i)$ indicates that the coefficient is specific to the block which the return series belongs to. Note that typically only few blocks are used so as to maintain the parsimony of the model.

2.3 Modelling the conditional mean

So far I have ignored any time regularities in the conditional means of returns. This may be justified when considering short-term, e.g. daily returns on financial assets. However, ignoring conditional mean dynamics becomes non-trivial when dealing with longer-term returns for commodities. Presence of autoregressive and/or moving-average components could be considered. A vector autoregressive moving average (vector ARMA, or VARMA) model could be employed to model them. Volatility models would then be used to describe the dynamics of the mean-unpredictable components of returns. That is, they would use residuals from conditional mean models rather than returns as dependent variables. I will assess the adequacy of conditional mean modelling by using Ljung-Box tests for autocorrelations.

When more than one candidate model is estimated for a given dependent variable, Bayesian information criterion (BIC) will be used to choose between them.

3 Data

I use weekly survey data on salmon spot prices from 1995 week 1 to 2013 week 13 (952 observations), provided by NOS Clearing ASA[3]. Prices are available for seven weight classes: 1-2 kg; 2-3 kg; 3-4 kg; 4-5 kg; 5-6 kg; 6-7 kg; and 7+ kg. Data are collected from exporters that purchase freshly slaughtered fish from salmon farmers. The following conditions apply for inclusion into the survey. Exporters must buy at least 2 truck-loads, or 40 metric tons (MT), of salmon every week and sell at least 50% of their total volume as gutted salmon to companies outside their own group. Purchases must be made from external farmers; this excludes within-company sales where both the exporting company and the farming company belong to the same owner. Prices must be agreed upon on the spot, thus no sales on forward contracts are included.

Around one-third of the total export volume is represented in the sample, but this share has varied from 8% to 47%. (The total production volume is close to the total export volume since almost all Norwegian farmed salmon is exported.) Survey results are published every Monday when the price data are around 10 days old: prices in the survey correspond to Thursday or Friday of the second-to-last week before the announcement day.

[3] The data may be downloaded at https://salmonprice.nasdaqomxtrader.com/historicalNOSprices.xls

As noted by one of the referees, the prices may be seen as auction prices similar to other commodity markets, e.g. the synchronous electricity implicit auction market. Also, they may be set non-synchronously within a given week. However, the possible non-synchronicity does not span across multiple weeks as the prices for all weight classes are set within a given week. Therefore, we may ignore it for the practical purposes of this study.

Weekly prices of selected weight classes are presented in Figure 1. The mean of volume-weighted average price in 1995-2013 was NOK 26.40/kg (figures in NOK are rounded to the first decimal) with a standard deviation of NOK 5.80/kg[4]. Price is increasing in fish size and the difference in the mean price of 1-2 kg salmon against 7+ kg salmon is NOK 6.00/kg.

Figure 1 Prices (NOK/kg) of selected weight classes

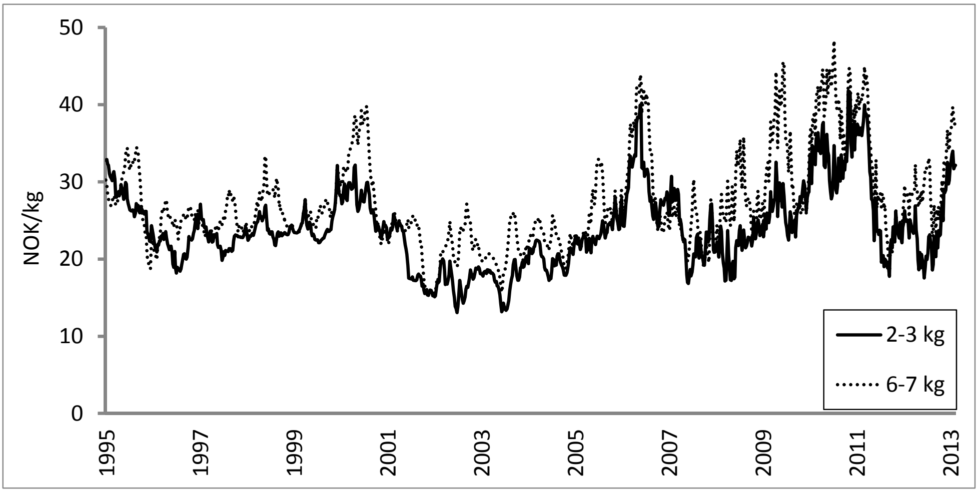

Logarithmic price returns represent a convenient transformation of prices and are often used in volatility analysis. Returns on salmon prices for all weight classes have means of 0.0% and standard deviations between 5.0% and 5.8%.

3.1 Structural change around 2006

There is a noticeable increase in variability of returns for all weight classes at around 2006. Apart from a couple of spikes in 1995, returns were mostly limited to $\pm 10\%$ until 2005, whereas since 2006 this threshold was frequently crossed. This is illustrated in Figure 2 for the 4-5 kg weight class, other weight classes showing similar patterns. Possible explanations for this change are discussed in Oglend (2013). The most prominent ones are the introduction of the maximum allowable biomass

[4] None of the figures is annualized here and in what follows.

(MAB) restriction in 2005; the opening of “Fish Pool” futures and options exchange in 2006; and the infectious salmon anemia (ISA) crisis in Chile and its aftermath in 2007-present (given the competition between the Norwegian and the Chilean salmon in the global market, the Chilean ISA crisis could be treated as a demand shock with respect to the Norwegian salmon).

Figure 2 Logarithmic returns on price for 4-5 kg fish

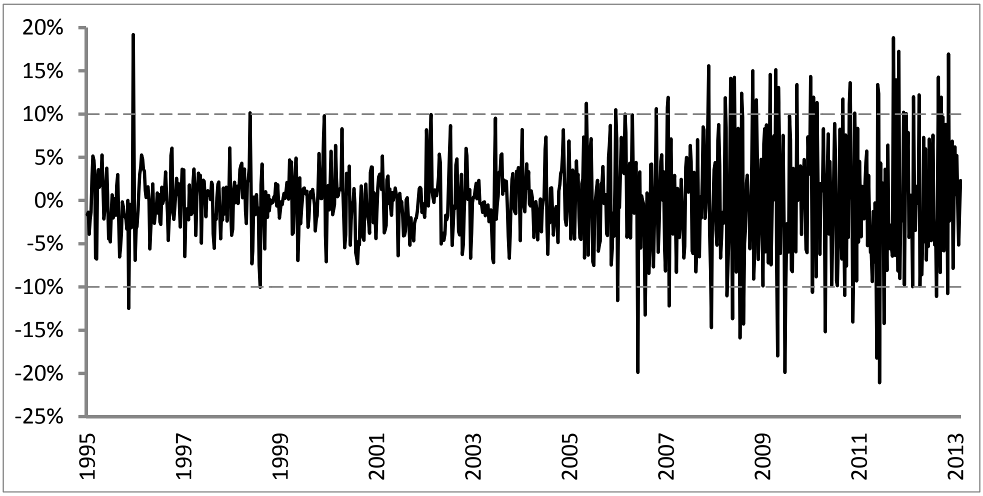

To formally test whether the variability of returns was different before and after 2006, I first need to formally separate the two periods. The basic methods for locating a change point in variance typically assume the data to be independent across time. This assumption may not be realistic in our case; however, accounting for the potential dependence explicitly would heavily complicate the analysis. Therefore, I ignore the dependence at this point, assuming that it is not strong enough to invalidate the results of the change point location procedure. Subsequent modelling results do not significantly violate this assumption.

The iterated cumulative sums of squares (ICSS) algorithm by Inclan and Tiao (1994) enables locating and formally establishing variance change points in a series of independent observations. Having applied ICSS for the seven weight classes separately, I find break points ranging from 2005 week 46 to 2007 week 27. I delete the period between these boundaries and obtain two new subsamples, 1996 week 1-2005 week 45 and 2007 week 27-2013 week 13. The year 1995 is omitted because of extraordinary events, namely, a forced cessation of feeding in salmon farms in the autumn (Asche & Guttormsen, 2009) and the subsequent introduction of feed quotas in 1996 (Aarset & Jakobsen, 2009).

Given the two subsamples, an F-test for the equality of variance in the two subpopulations could be used. However, the F-test is sensitive to departures from normality of the data, which are revealed below. Therefore, Levene's test (Levene, 1960) or Brown-Forsythe test (Brown & Forsythe, 1974) should be used instead. The results of the two tests applied separately on every weight class strongly reject the equality of variances with p-values all below 0.01. Strictly speaking, the latter results are conditional on the data-dependent split of the original sample into the two subsamples. However, the differences in variances seem large enough so that this point of concern may be ignored.

Due to the unequal variances of returns between the former and the latter period, modelling conditional means and conditional variances through the whole sample is problematic. In the subsequent analysis I therefore distinguish between 1996 week 1-2005 week 45, or the first half (H1) of the original sample, and 2007 week 27-2013 week 13, or the second half (H2) of the original sample.

3.2 Seasonality

There are seasonal patterns in salmon prices; relative prices across weight classes vary systematically in one-year cycles (Asche & Guttormsen, 2001). Thus, large fish occasionally becomes cheaper than small fish. This is evident from Figure 3.

The seasonality is due to the nature of salmon production with slow fish growth in cold periods (January-April) and more rapid growth when the sea is warmer (July-October) (Marine Harvest, 2013). Seasonality in production volume is illustrated in Figure 4 which depicts volume distribution across selected weight classes.

The relative volume shares of different weight classes have changed considerably over the sample period. The share of smaller fish has increased while the share of larger fish has decreased. Also, the annual seasonal patterns may have evolved over the sample period e.g. due to the major industry events as listed before.

Ideally, seasonal adjustment would be done simultaneously with the econometric modelling of the time series. However, this is not feasible in our case due to a high computational burden. Therefore, I perform seasonal adjustment as a separate, initial stage of the modelling. To adjust for the seasonality, I consider each subsample, each series separately. I use a regression with ARMA errors, with Fourier terms and an Easter dummy as regressors. The set of Fourier terms is made of pairs of $sin(\cdot)$ and $cos(\cdot)$ series with periodicity of 1 year, ½ year, ¼ year etc. The ARMA order for the errors is allowed to be any subset of an ARMA(4,4) specification. The number of pairs of Fourier terms, the inclusion or omission of the Easter dummy and the ARMA specification are selected simultaneously using the corrected Akaike information criterion (AICc). For H1, three to eight pairs of

Fourier terms are selected, while one to five pairs suffice for H2. The use of AICc allows for more flexibility than using BIC which could otherwise be preferred due to its consistency. (BIC selects only one pair of Fourier terms in seven out of eight cases in H1 and six out of eight cases in H2; this may be overly restrictive and may result in a part of the seasonal effect still remaining after the adjustment.) The sum of the effects of the Fourier terms and the Easter dummy constitutes the seasonal component, which is removed from the data before proceeding. From this point forward, seasonally-adjusted series will be used in place of the original series without explicit reference.

Figure 3 Prices (NOK/kg) of selected weight classes minus volume-weighted average price

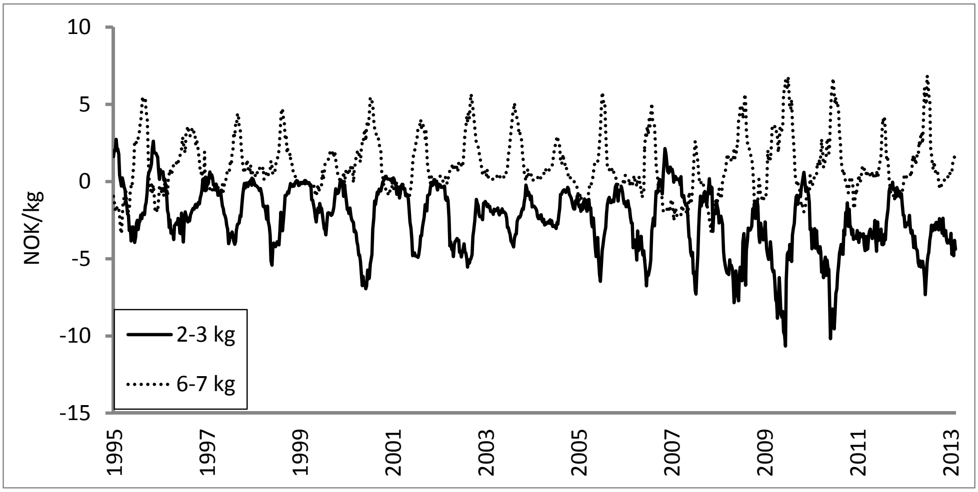

Figure 4 Relative volume shares of selected weight classes

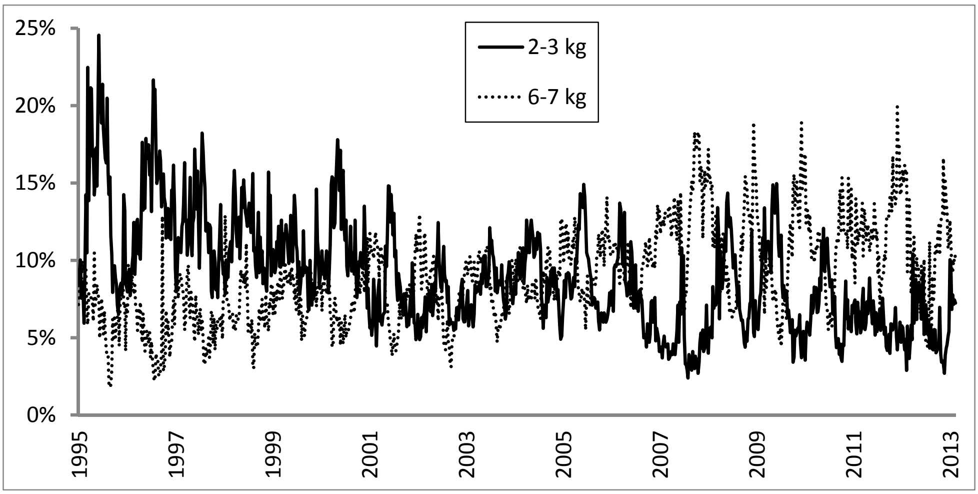

3.3 Descriptive statistics

Distributions of returns for the seven weight classes over 1996-2005 are centered at zero with standard deviations between 3.1%-3.7% (Table 1). The distributions are not normal as evidenced by Shapiro-Francia test results, with excess kurtosis being positive and significantly different from zero by the Anscombe-Glynn test. That is, returns are more densely concentrated in a narrow interval around the mean and also more spread out in tails, but less in between, as compared with the normal distribution. Meanwhile, skewness is positive and statistically different from zero by the d'Agostino test for all but the smallest and the largest fish; thus large positive returns are more common than large negative returns. The largest negative returns are -8% to -15% and the largest positive returns are +11% to +17%. The 1-2 kg weight class has the largest negative and the largest positive single return. Since this weight class has the smallest volumes, its extensive variability may come from the relatively small samples in the NOS weekly survey over the period analyzed.

Table 1 Descriptive statistics of logarithmic returns on seasonally-adjusted prices, 1996-2005

Weight class	Mean x100	Median x100	St. dev. x100	Skewness	Excess kurtosis	Shapiro-Francia test	Min. x100	Max. x100	Leverage effect†
1-2 kg	-0.03	-0.03	3.67	0.06	2.02***	0.9744***	-15.33	16.54	0.06
2-3 kg	0.00	-0.05	3.09	0.29*	0.97***	0.9844***	-10.12	11.46	0.20***
3-4 kg	0.01	-0.09	3.13	0.39**	1.20***	0.9818***	-11.26	13.50	0.25***
4-5 kg	0.01	-0.03	3.13	0.33*	0.62**	0.9898***	-8.61	11.81	0.23***
5-6 kg	0.01	-0.17	3.19	0.30*	0.53**	0.9914***	-8.36	11.43	0.22***
6-7 kg	0.01	-0.37	3.12	0.34*	0.76***	0.9880***	-9.91	11.32	0.23***
7+ kg	-0.01	0.12	3.42	0.27	1.49***	0.9793***	-11.87	13.08	0.18***
Avg.	0.01	-0.17	2.96	0.40**	0.74***	0.9862***	-8.03	11.26	0.27***

Note: ***, ** and * mark statistical significance at 1%, 5% and 10% level, respectively. Significance applies to mean, skewness, excess kurtosis, Shapiro-Francia test and leverage effect.

† Correlation between squared returns and prices.

Zero mean of weekly returns cannot be rejected for the period of 2007-2013 (see Table 2). Meanwhile, standard deviations roughly double from H1 and range between 6.6%-7.5% in H2. Contrary to the results for H1, normality of return distributions in H2 cannot be rejected using the Shapiro-Francia, d'Agostino and Anscombe-Glynn tests, except for the 1-2 kg fish (probably due to the small samples' problem mentioned above). Extreme returns are larger in H2 than in H1: -19% to -27% and +18% to +22% where the 1-2 kg weight class again has the most extreme values.

Table 2 Descriptive statistics of logarithmic returns on seasonally-adjusted prices, 2007-2013

Weight class	Mean x100	Median x100	St. dev. x100	Skewness	Excess kurtosis	Shapiro-Francia test	Min. x100	Max. x100	Leverage effect
1-2 kg	0.17	-0.14	6.61	-0.01	1.31***	0.9807***	-26.51	22.07	0.08
2-3 kg	0.17	-0.38	7.15	0.20	0.01	0.9930	-18.51	20.65	0.23***
3-4 kg	0.13	-0.30	7.34	0.10	-0.24	0.9945	-20.04	19.09	0.15***
4-5 kg	0.14	-0.16	7.23	0.05	-0.24	0.9949	-20.73	17.78	0.12***
5-6 kg	0.16	-0.24	7.33	0.08	-0.21	0.9948	-21.22	19.15	0.14***
6-7 kg	0.21	-0.07	7.42	0.14	-0.16	0.9922	-22.84	19.22	0.19***
7+ kg	0.25	-0.17	7.51	0.13	-0.20	0.9941	-23.45	18.77	0.20***
Avg.	0.15	-0.32	7.28	0.12	-0.26	0.9928	-20.34	17.94	0.17***

Note: see note under Table 1.

I document not only a quantitative change, i.e. a large increase in the standard deviations, but also a qualitative change, i.e. emergence of normality, when moving from H1 to H2. Had I not recognized this shift and proceeded by analyzing the entire sample, the results would have been qualitatively and quantitatively different. For example, I would have found very pronounced indications of non-normality based on the Shapiro-Francia and Anscombe-Glynn tests.

3.4 Empirical evidence on stylized facts

Recalling the stylized facts of asset returns from Cont (2001), I test for gain-loss asymmetry. It is present in H1 (except for the smallest and the largest fish) since skewness was found to be statistically different from zero. Meanwhile, it is absent in H2. Unconditional heavy tails are found in H1 but not in H2 (except for the 1-2 kg weight class), based on the values of excess kurtosis and the Anscombe-Glynn test results.

Correlation between price level and conditional variance, or leverage effect, is expected to be positive as noted in Oglend (2013) and references therein. Starting with the full sample (1995-2013) and using squared returns as proxies for the conditional variances, I find correlation to be close to zero for the smallest fish but moderately positive (up to 0.15) and statistically significant for all other weight classes and the average price. This pattern remains when I divide the sample into the two subsamples, with the highest correlations reaching 0.25 in H1 (0.27 for returns on the average price) and 0.23 in H2. Tables 1 and 2 report the details.

Correlation between trading volume and return volatility is difficult or even impossible to assess for salmon since I do not have exact trading volume data, nor a good proxy for it.

The Ljung-Box test indicates statistically significant autocorrelations for all weight classes in both subsamples. Next, I test for presence of autocorrelations in absolute values of returns. I find them statistically significant for all weight classes in H1. On the other hand, only weak evidence of autocorrelations in absolute values may be detected in H2, except for the smallest fish. However, the Ljung-Box test statistic does not have its usual null distribution in presence of conditional heteroskedasticity (Wooldridge, 1991). Therefore, I will only be able to interpret the Ljung-Box test results properly after testing for conditional heteroskedasticity, which I do next.

Oglend and Sikveland (2008) and Oglend (2013) found volatility clustering in the returns on the average salmon prices. Here I ask whether their conclusions hold for weight-class-specific prices in the separate subsamples H1 and H2. Using an ARCH-LM test on returns in H1, I find evidence of conditional heteroskedasticity (which invalidates the Ljung-Box test results above) for all weight classes but the smallest fish. Meanwhile, ARCH effects in H2 are marginal. These conclusions are robust to varying the number of lags in the ARCH-LM test. Thus we see a qualitative change from H1 to H2. Volatility clustering in returns was present until 2005 but almost vanished after 2006. This observation would have been missed had I tested over the entire period 1995-2013, where the ARCH-LM test clearly indicates presence of ARCH effects.

All return series display autocorrelations and some of them also have conditional heteroskedasticity. I therefore proceed by explicitly modelling these features in the next section. Before that, I document that the characteristics of returns on salmon prices only partly match the stylized facts of asset returns. Namely, leverage effect is the single stylized fact from Cont (2001) that is present in the majority of weight classes throughout both subsamples H1 and H2. This exemplifies the difference between salmon, a partly-storable commodity, and pure financial assets.

4 Empirical results

To model the autocorrelations in returns, I could choose from a class of VARMA models with returns on the EUR/NOK rate as an exogenous variable. Although salmon trade may in principle have an impact on the EUR/NOK rate, it is likely to be negligibly small, hence the exogeneity assumption for the EUR/NOK rate. However, the dimension of the data is high relative to the sample size: seven series of weight-class-specific returns and the return series for the EUR/NOK rate against 463 observations in H1 and 298 observations in H2. Thus the “curse of dimensionality” is an issue. To avoid it, I restrict the analysis to the class of univariate autoregressive moving average models, ARMA(p,q), where p is the autoregressive order and q is the moving-average order. I also include returns on the EUR/NOK rate as an exogenous regressor. I estimate all models nested in ARMA(4,4) for each weight class separately. BIC is then used to select the optimal model from the pool of all estimated models for a given weight class.

4.1 Univariate models for 1996-2005

To evaluate the relevance of the BIC-selected models for the subsample H1, I examine model residuals and test them for presence of serial correlation and conditional heteroskedasticity. The model for the 1-2 kg weight class has satisfactory residual properties. It is an MA(3) model with MA(1) and MA(2) terms restricted to zero. For all other weight classes, ARCH patterns are present in the residuals, which suggests adding a GARCH(m,s) structure to the models.

The following procedure is applied for each weight class. I try orders m from 0 to 3 and s from 0 to 1 where $s > 0$ only if $m > 0$. That is, I estimate all models nested in ARMA(3,3)-GARCH(3,1) with the GARCH term allowed to be present only if at least one of the ARCH terms is present. The maximum ARMA order is reduced from (4,4) to (3,3) to lessen the computational burden to an acceptable level. Again, BIC is used to select one model from the pool of all estimated models.

This approach has proved fruitful, and I now find models with well-behaved residuals for each weight class and also for the returns on the average price. Lower BIC values suggest that the new models of the form ARMA(p,q)-GARCH(m,s) are to be preferred over the initial ARMA(p,q) models. Model coefficients and their robust standard errors are presented in Table 3. For models producing standardized residuals that follow the assumed theoretical distribution (in our case, the normal distribution), non-robust standard errors would be preferred. However, I remain conservative and report robust standard errors everywhere, also for direct comparability across models.

AR(2)-ARCH(1) is the selected specification for five out of six weight classes and also for the returns on the average price. The single exceptional case is ARMA(2,3)-ARCH(2) specification for the 2-3 kg weight class. These are the smallest fish among the weight classes considered at this stage, and their supply is somewhat scarce. Therefore, it may be natural that their return dynamics differs from the more abundant weight classes of moderate-size fish. Returns on the EUR/NOK rate are not selected as explanatory variables for any of the weight classes. Thus I infer that fluctuations in the EUR/NOK rate were a relatively minor determinant (if at all) of fluctuations in salmon prices in H1.

Although the models are very parsimonious, their R^2 values reach 0.19-0.23 for the 3-6 kg weight classes and also for the returns on the average price. R^2 is above 0.10 for the 2-3 kg and the 6-7 kg but negligibly small for the extreme-sized fish. The relatively high R^2 values for the most abundant weight classes are somewhat surprising as return dynamics is not supposed to be easily explainable beyond the seasonal patterns. If it was, market forces from both the supply side and the demand side could work to exploit the regularities. On one hand, it could be profitable for salmon farmers to change their supply schedule so as to capitalize on periods of higher prices. On the other hand, salmon buyers could adjust their purchasing schedule to acquire the fish when they are cheaper. Thus one may hypothesize that either the market participants did not have the physical

Table 3 Model coefficients and robust standard errors (in parentheses), 1996-2005

Weight class	Intercept x100	AR1	AR2	AR3	MA1	MA2	MA3	EUR/NOK	ω x100^2	ARCH1	ARCH2	ARCH3	GARCH1	EUR/NOK
1-2 kg	-0.0348	---	---	---	---	---	-0.1343	---	---	---	---	---	---	---
	(0.1464)						(0.0475)							
2-3 kg	-0.0562	0.3476	-0.1668	---	---	---	-0.1060	---	5.5051	0.2166	0.1334†	---	---	---
	(0.1326)	(0.0568)	(0.0526)				(0.0583)		(0.7685)	(0.0827)	(0.1042)			
3-4 kg	-0.1595	0.4492	-0.3684	---	---	---	---	---	5.8752	0.2586	---	---	---	---
	(0.1466)	(0.0463)	(0.0482)						(0.7274)	(0.0996)				
4-5 kg	-0.1212	0.4652	-0.3429	---	---	---	---	---	5.5962	0.2428	---	---	---	---
	(0.1417)	(0.0478)	(0.0430)						(0.5676)	(0.0705)				
5-6 kg	-0.1265	0.4330	-0.2834	---	---	---	---	---	6.3914	0.2199	---	---	---	---
	(0.1614)	(0.0527)	(0.0454)						(0.6716)	(0.0891)				
6-7 kg	-0.1989	0.4469	-0.2336	---	---	---	---	---	5.0698	0.4469	---	---	---	---
	(0.1698)	(0.0434)	(0.0468)						(0.7099)	(0.1035)				
7+ kg	-0.1033	0.1795	-0.0937†	---	---	---	---	---	8.2647	0.2975	---	---	---	---
	(0.1690)	(0.0594)	(0.0583)						(1.1634)	(0.1117)				
Avg.	-0.1388	0.4628	-0.3491	---	---	---	---	----	4.7908	0.3152	---	---	---	---
	(0.1414)	(0.0410)	(0.0398)						(0.5392)	(0.0769)				

Note: † not significant at 10% significance level; all other coefficients, except for intercepts, are statistically significant at 10% or lower significance level.

Table 4 Model diagnostics, 1996-2005

Weight class	R^2	Fit of cond. variance	Shapiro-Francia test	Skewness of standardized residuals	Excess kurt. of standardized residuals	Ljung-Box (d.f.+1 lag) test	Ljung-Box (d.f.+5 lags) test	Li-Mak (d.f.+1 lag) test	Li-Mak (d.f.+5 lags) test
1-2 kg	0.0171	---	0.9745***	0.06	2.12***	3.26*	8.11	0.62†	4.05†
2-3 kg	0.1062	0.0365	0.9887***	0.20	0.77***	0.57	3.19	0.37	10.39*
3-4 kg	0.2031	0.0315	0.9891***	0.20	0.78***	0.93	5.03	2.38	6.86
4-5 kg	0.2311	0.0606	0.9957	0.14	0.29	0.22	7.70	0.01	3.54
5-6 kg	0.1898	0.0185	0.9960	0.03	0.33	0.38	4.81	0.25	0.64
6-7 kg	0.1469	0.0362	0.9962	0.15	0.16	0.93	4.28	1.58	2.26
7+ kg	0.0137	0.0308	0.9819***	0.29*	1.48***	1.81	7.71	0.54	4.22
Avg.	0.2127	0.0637	0.9936**	0.26	0.29	0.26	7.30	2.25	5.46

Note: fit of cond. variance is measured as squared correlation between squared residuals from the conditional mean model and the estimated conditional variances. Shapiro-Francia test is applied on standardized residuals. Ljung-Box test statistic does not have the usual null distribution when applied on standardized residuals of a GARCH model; still, I report the test statistic for completeness. d.f. stands for the number of parameters in conditional mean model, excluding intercept, for the Ljung-Box test. d.f. stands for the number of parameters in the conditional variance model for the Li-Mak test. ***, ** and * mark statistical significance at 1%, 5% and 10% level, respectively.

† ARCH-LM(2) and ARCH-LM(5) tests were used in place of Li-Mak(d.f.+1) and Li-Mak (d.f.+5) tests, respectively, for the model without the conditional variance part.

capability of exploiting the price patterns (e.g. due to limited throughput of salmon processing plants or other technological constraints) or that they did not use suitable econometric models to spot the regularities in the conditional mean of returns. However, I have more data at hand than were available at the time of actual decision making, which weakens the latter hypothesis. Also, statistical significance need not imply economic significance; that is, the size of the potential economic gain due to an effective utilization of the return regularities is not trivial.

The estimated shocks to returns, i.e. the non-standardized model residuals, are non-normal for all weight classes and for the returns on the average price. They have positive excess kurtosis, while skewness is largely absent. Hence shocks to returns are symmetric but have more extreme values than a normally distributed random variable. Explicit model diagnostics may be found in Table 4.

In the conditional variance models, statistically significant and positive ARCH(1) terms suggest that the size of the current shock tends to affect the size of the next period's shock. The same holds true for the second-to-next period due to the positive ARCH(2) term for the 2-3 kg weight class. This is enough to create a cascade effect: a large shock in the current period makes it likely that a large shock will follow in the next period. This holds recursively for every period in the future. The cascade is however diminishing as time passes. The rate of diminution may be quantified by the sum of ARCH coefficients, i.e. the persistence parameter λ, which ranges between 0.22 and 0.45 in H1. These values are smaller than 0.5, thus I conclude that one period is enough for the volatility to move more than halfway towards its long-term mean. This is a rather fast rate and may be due to weekly data being a somewhat low observation frequency in the context of GARCH model applications. ω expressed as a share of the unconditional variance ranges from 0.52 to 0.77. Thus a half to three-quarters of the unconditional variance is the basis for conditional variance; shocks bring the conditional variance up from the basis, but it would come down shortly after, if not for new shocks. The absence of GARCH terms, or autoregressive effects, in the conditional variance equations implies finite memory in the conditional variance process and confirms that the effects of shocks are relatively short lived. The squared returns on the EUR/NOK rate are not selected as regressors for any weight class. Hence, fluctuations in the exchange rate do not seem to add to the volatility of salmon prices, which is in line with the results in Oglend (2013).

To assess how well the models capture the patterns of conditional variance, I calculate squared correlations between squared, non-standardized model residuals and the fitted values of the conditional variance. The squared correlations range from 0.02 to 0.06. At first glance this looks rather low, suggesting that only a small part of the variation in the conditional variances is of ARCH type. However, we do not actually observe the conditional variances; the squared residuals are only proxies for them. Also, (G)ARCH models imply a perfect fit for the conditional variance by construction. Therefore, I remain careful and treat the squared correlations merely as indirect indicators of model fit. (See Andersen and Bollerslev (1998) for a thorough treatment of the matter.) Unfortunately, I cannot compare my results with other studies as it is uncommon to report the fit of conditional variance models.

4.2 Univariate models for 2007-2013

Models for H2 are presented in Table 5. A high degree of similarity across weight classes is evident; all the models have the same order, only the magnitudes and the significance of estimated coefficients vary. The conditional means of returns are explained by MA(1) and MA(2) terms in place of the autoregressive terms as of H1. Returns on the EUR/NOK rate are included for all weight classes and also for the returns on the average price. Although the EUR/NOK coefficients are significant for only two out of seven weight classes and in the model for the returns on the average price, apparently the variable has considerable explanatory power; otherwise it would not have been selected using BIC. The coefficients have the expected positive sign such that depreciation in NOK against EUR leads to an increase in salmon prices. Coefficient values are between 0.34-0.38 for the small fish (1-3 kg) and between 0.47 and 0.64 for all the other weight classes and for the returns on the average price. Thus a 1% depreciation in NOK against EUR yields between 0.34% and 0.64% increase in the salmon price for each weight class.

Table 6 provides model diagnostics for H2. R^2 values are relatively low at 0.10-0.15 with the lowest values recorded for the least abundant weight classes of 1-3 kg and 7+ kg. Compared with H1, R^2 values have decreased, except for the smallest and the largest fish. The lower R^2 values, together with the higher unconditional variance of returns, imply that the ARMA patterns are more lost in noise in H2 than in H1. The increases in R^2 for the smallest and the largest fish may be due to the inclusion of the returns on the EUR/NOK rate as an explanatory variable. Since this is a contemporaneous regressor, the models are no longer directly applicable to prediction.

The loss of explanatory power of the ARMA models may have several explanations. First, the market participants may have detected the patterns in returns and adjusted their operations accordingly, eliminating the regularities. Second, the introduction of the futures contracts by “Fish Pool” and the resulting availability of timely price information may have affected the price discovery process. Also, a change in salmon supply schedule due to the MAB restriction since 2005 may have played a role.

Normality of the (non-standardized) shock distributions cannot be rejected for the returns on the average price and for all weight classes with the exception of the 1-2 kg fish. This is in contrast to the earlier period, where non-normality was present due to positive excess kurtosis. Thus the extreme values abundant in H1 became less common in H2, when assessed relative to the means and variances of the (non-standardized) shock distributions. Meanwhile, standardized shocks are all normally distributed, except for the smallest fish.

All the weight classes share qualitatively similar volatility dynamics. A single ARCH(3) term is included in the models for the conditional variance. The coefficient values range around 0.10-0.12 for the 3+ kg weight classes and the returns on the average price, while values between 0.19 and 0.29 are estimated for the 1-3 kg weight classes. The economic interpretation of the ARCH(3) pattern is somewhat unintuitive. A current shock tends to have an impact on volatility three periods later, but not before. Since the coefficients are small, I may claim that the ARCH patterns are weak in H2. This falls in line with the weaker signal-to-noise ratio in the models for the conditional mean. The half-life parameter K is well below one for all weight classes, indicating that the effects of shocks are short lived. The squared returns on the EUR/NOK rate are not included in the conditional variance equations; this is in agreement with Oglend (2013). Therefore, the effect of the EUR/NOK rate on salmon price is limited to the conditional mean.

Regarding the fit of the conditional variance models, squared correlations between the squared residuals and the fitted conditional variances remain low at 0.03-0.05 in H2. Thus the increase in the variance of shocks from H1 to H2 may be mostly not of GARCH type, and the clustering of small and large returns becomes less pronounced by 2007-2013. This is also echoed by BIC preferring models with constant conditional variance (i.e., models without the conditional variance part) for three out of seven weight classes and for the returns on the average price. However, models without the conditional variance part show statistically significant volatility clustering in the residuals. Since the focus of this study is on volatility, I retain the conditional variance models even though BIC suggests they are superfluous.

So far I have explored the univariate time regularities of returns for each weight class separately. Links between the mean-unpredictable components of individual series of returns will now be analyzed using DCC models.

4.3 DCC modelling

I model the conditional covariances separately in H1 and H2. In each of the subsamples, I specify an asymmetric DCC(3,3) model and also all models nested in it. Clearly, the standard (symmetric) DCC models are included in this set. I also specify a number of flexible DCC(1,1) models with two and three blocks. (For flexible DCC models, lag order greater than (1,1) is not available in the R software (Ghalanos, 2014).) The split into blocks is monotonous in fish size. I consider all possible block structures of two and three blocks where all weight classes in each block are of consecutive weights. Examples of the permitted block structure are the following: 1-2 kg and 2+ kg; 1-5 kg and 5+ kg; 1-2 kg, 2-3 kg and 3+ kg; 1-5 kg, 5-7 kg and 7+ kg. On the other hand, the following block structure is

Table 5 Model coefficients and robust standard errors (in parentheses), 2007-2013

Weight class	Intercept x100	AR1	AR2	AR3	MA1	MA2	MA3	EUR/NOK	ω x100^2	ARCH1	ARCH2	ARCH3	GARCH1	EUR/NOK
1-2 kg	0.1530	---	---	---	-0.2661	-0.1962	---	0.3835†	28.6110	---	---	0.2885	---	---
	(0.1753)				(0.0529)	(0.0560)		(0.2845)	(3.8844)			(0.1330)		
2-3 kg	0.1427	---	---	---	-0.1817	-0.3056	---	0.3393†	36.9729	---	---	0.1946	---	---
	(0.1839)				(0.0465)	(0.0468)		(0.2735)	(4.4770)			(0.0805)		
3-4 kg	0.1411	---	---	---	-0.1916	-0.3469	---	0.6402	41.0980	---	---	0.1033	---	---
	(0.1725)				(0.0506)	(0.0483)		(0.3075)	(3.9686)			(0.0539)		
4-5 kg	0.1542	---	---	---	-0.2075	-0.3410	---	0.6194	39.8135	---	---	0.0981	---	---
	(0.1661)				(0.0499)	(0.0462)		(0.3077)	(3.7048)			(0.0529)		
5-6 kg	0.1238	---	---	---	-0.2206	-0.3396	---	0.4683†	39.6850	---	---	0.1238	---	---
	(0.1599)				(0.0493)	(0.0431)		(0.3196)	(3.7358)			(0.0563)		
6-7 kg	0.1973	---	---	---	-0.1675	-0.3279	---	0.5978†	42.3878	---	---	0.1099	---	---
	(0.1899)				(0.0510)	(0.0448)		(0.3801)	(3.9638)			(0.0606)		
7+ kg	0.1782	---	---	---	-0.1828	-0.3399	---	0.6059†	41.7203	---	---	0.1217	---	---
	(0.1780)				(0.0480)	(0.0416)		(0.3900)	(3.9658)			(0.0678)		
Avg.	0.1596	---	---	---	-0.2051	-0.3417	---	0.6036	40.3741	---	---	0.0977	---	---
	(0.1671)				(0.0491)	(0.0447)		(0.3118)	(3.6859)			(0.0519)		

Note: † not significant at 10% significance level; see also note under Table 3.

Table 6 Model diagnostics, 2007-2013

Weight class	R^2	Fit of cond. variance	Shapiro-Francia test	Skewness of standardized residuals	Excess kurt. of standardized residuals	Ljung-Box (d.f.+1 lag) test	Ljung-Box (d.f.+5 lags) test	Li-Mak (d.f.+1 lag) test	Li-Mak (d.f.+5 lags) test
1-2 kg	0.1008	0.0370	0.9884**	0.13	0.77**	1.15	5.05	0.20	2.88
2-3 kg	0.1000	0.0378	0.9939	0.24	0.01	0.74	7.33	0.01	2.60
3-4 kg	0.1476	0.0279	0.9944	0.21	-0.20	0.10	4.10	0.04	1.07
4-5 kg	0.1544	0.0323	0.9957	0.14	-0.30	0.19	4.46	0.26	1.09
5-6 kg	0.1493	0.0471	0.9961	0.12	-0.36	0.46	7.21	0.00	0.54
6-7 kg	0.1286	0.0319	0.9950	0.12	-0.42*	1.16	4.31	0.28	1.30
7+ kg	0.1102	0.0314	0.9938	0.11	-0.48**	2.79*	5.44	0.73	3.28
Avg.	0.1535	0.0309	0.9936	0.21	-0.34	0.35	5.28	0.05	0.45

Note: see note under Table 4.

not permitted: any structure where one of the blocks consists of the 1-2 kg and the 3-4 kg weight classes only; this is not permitted due to the non-consecutive weights as the 2-3 kg weight class is excluded from the block. I estimate all the asymmetric DCC and flexible DCC models and select only one of them using BIC.

4.3.1 DCC model for 1996-2005

BIC selects a very parsimonious model for H1:

$$q_{ij,t|t-1} = (1 - 0.18 - 0.07)\rho_{ij} + 0.18\eta_{i,t-1}\eta_{j,t-1} + 0.07neg(\eta_{i,t-3})neg(\eta_{j,t-3})$$

The moving average coefficient γ_1 equals 0.18; hence, the cross product of the standardized most recent shocks has a relatively small effect on the conditional correlations. The asymmetry-governing coefficient θ_3 is 0.07, suggesting a mild effect of negative shocks lagged by three periods. The unconditional correlation has the largest weight between the three components, amounting to 0.74 (not 0.75, due to rounding), which shows that the conditional correlations are relatively close to their unconditional counterparts. The half-life parameter *K* equals 0.3; thus more than a half of the gap between a given conditional correlation and its long-term mean (the unconditional correlation) is closed in one period.

Figure 5 depicts the estimated conditional correlations for all pairs of weight classes in both H1 and H2. The conditional correlations are positive throughout H1 with rare exceptions. The mean-unpredictable components of returns for salmon of similar size, or from neighboring weight classes, are more correlated than for salmon of dissimilar size, or from distant weight classes. Average conditional correlations, which approximately equal the unconditional correlations, range from 0.6 to 0.9 between the neighboring weight classes. They are highest for the pairs of the most abundant weight classes, i.e. the 3-4 kg with 4-5 kg pair and the 4-5 kg with 5-6 kg pair. The smallest average conditional correlations of around 0.3 to 0.5 are observed between the 2-3 kg and 7+ kg weight classes and in pairs where the 1-2 kg weight class is involved. The latter may be explained by the scarcity and the small samples of the 1-2 kg and the 7+ kg weight classes. See Table 7 for the complete set of the unconditional correlations in H1.

Looking at the time development of the conditional correlations through 1996-2005, they are relatively stable between the neighboring weight classes. This is especially characteristic to the pairs of the most abundant weight classes: conditional correlations for the 3-4 kg with 4-5 kg pair and the 4-5 kg with 5-6 kg pair remain above 0.8 for 97% of the time in H1 and have standard deviations of only 0.04-0.05. This may result from common market trends and also the similarity of the one-time shocks for fish of similar size. For example, salmon of the neighboring weight classes will naturally be easiest to substitute between when a demand shock occurs. Meanwhile, the conditional correlations vary more for distant weight classes, e.g. the standard deviation of the conditional correlations

Figure 5 Conditional correlations between shocks to logarithmic price returns for pairs of weight classes

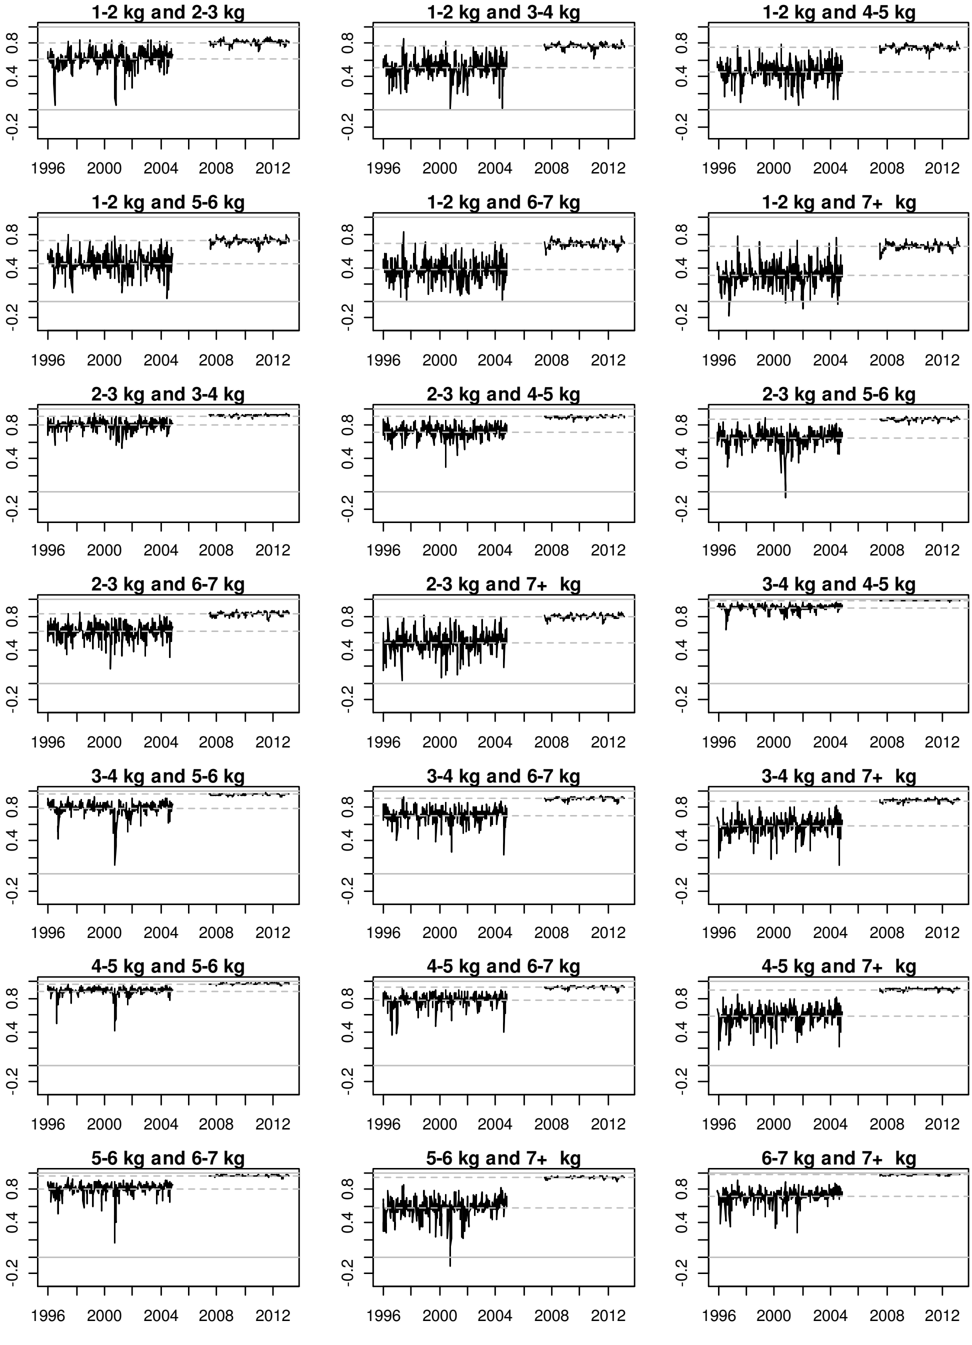

Note: solid horizontal lines mark zero and one; dashed horizontal lines mark unconditional correlations for 1996-2005 and 2007-2013.

between the 1-2 kg and the 7+ kg weight classes is 0.13. This may be due to the dissimilarity of shocks affecting distant weight classes. When a disease hits a salmon farming site, all fish in that site may be slaughtered and dumped on the market. If the fish are of uniform size at the site, the supply shock will be concentrated on one weight class but uncorrelated with other weight classes.

I see little evidence of any deterministic time trend in the conditional correlations in H1, although this may be due to a lack of the relevant flexibility in the DCC model.

It is worth noting that BIC values of the best few DCC models in H1 are rather concentrated; none of the models clearly dominates the others. One of the closest competitors to the model discussed above is a flexible DCC(1,1) model with two blocks of weight classes, 1-3 kg and 3+ kg. This block structure suggests that the "outlier" fish in terms of conditional correlations dynamics are the two smallest weight classes, as opposed to, for example, the few largest weight classes. On the other hand, the difference between the two blocks is relatively small as ultimately neither this model nor any of the other flexible DCC models is selected by BIC.

Table 7 Unconditional correlations between estimated shocks to returns, 1996-2005

Weight class	1-2 kg	2-3 kg	3-4 kg	4-5 kg	5-6 kg	6-7 kg	7+ kg
1-2 kg	1.00	0.57	0.49	0.43	0.43	0.35	0.30
2-3 kg		1.00	0.79	0.72	0.64	0.61	0.48
3-4 kg			1.00	0.90	0.77	0.70	0.57
4-5 kg				1.00	0.89	0.79	0.59
5-6 kg					1.00	0.81	0.57
6-7 kg						1.00	0.70
7+ kg							1.00

4.3.2 DCC model for 2007-2013

The BIC-selected model for H2 is the classic DCC(1,1):

$$q_{ij,t|t-1} = (1 - 0.06 - 0.72)\rho_{ij} + 0.06\eta_{i,t-1}\eta_{j,t-1} + 0.72q_{ij,t-1|t-2}$$

The unconditional correlation no longer has a large weight, $1 - \gamma - \delta = 0.22$. The moving-average effect $\gamma = 0.06$ is fairly small in absolute terms and smaller than in H1. On the other hand, the autoregressive memory, or inertia, $\delta = 0.72$, is present and strong. That is, the conditional correlation would remain off its long-term mean for a longer period of time and would only slowly close the gap. The half-life parameter $K = 2.7$; hence, it takes almost three periods to close half of the gap between a given conditional correlation and its long-term mean.

Analysis of different pairs of weight classes (see Figure 5) yields similar conclusions to those drawn for H1. The neighboring weight classes show higher conditional correlations while the distant weight classes have lower ones. However, we see pronounced increases in correlations for all pairs of weight classes when moving from H1 to H2. The uniformity of the increase across pairs may partly be attributed to the inflexibility of the DCC model. Still, experimenting with smaller, pairwise DCC models yields qualitatively similar results. This confirms that shocks to returns became more correlated across weight classes in the latter period. Even the distant weight classes show average conditional correlations larger than 0.6, while average conditional correlations between the neighboring weight classes all exceed 0.9, except for the 1-2 kg with 2-3 kg pair. The linear association between shocks for the most abundant weight classes, the 3-4 kg with the 4-5 kg and the 4-5 kg with the 5-6 kg, is essentially perfect. Conditional correlations of the first pair never fall below 0.95 and for the second pair they remain above 0.95 for 99% of the time through 2007-2013. We also see that with the overall upward shift of the conditional correlations from H1 to H2, the differences in the conditional correlations between different pairs of weight classes have diminished. The complete set of unconditional correlations in H2 is provided in Table 8.

Table 8 Unconditional correlations between estimated shocks to returns, 2007-2013

Weight class	1-2 kg	2-3 kg	3-4 kg	4-5 kg	5-6 kg	6-7 kg	7+ kg
1-2 kg	1.00	0.81	0.76	0.75	0.72	0.69	0.67
2-3 kg		1.00	0.92	0.90	0.87	0.83	0.81
3-4 kg			1.00	0.99	0.96	0.91	0.89
4-5 kg				1.00	0.98	0.94	0.92
5-6 kg					1.00	0.97	0.95
6-7 kg						1.00	0.99
7+ kg							1.00

As regards time patterns, amplitudes of the conditional correlations are substantially smaller in H2 as compared with H1. The largest standard deviation of the conditional correlations is 0.04 for the pair of the smallest against the largest fish; meanwhile, the standard deviations are below 0.01 for both the 3-4 kg with 4-5 kg pair and the 4-5 kg with 5-6 kg pair of weight classes. The whole range of conditional correlations in H2 is above the average conditional correlations (marked by dashed horizontal lines in Figure 5) in H1 for all pairs. Thus not only have shocks become more correlated across weight classes on average, but the correlations have also turned more stable with respect to time in H2. In a way, salmon has become a more homogenous product with respect to fish size.

Similarly to H1, the few best DCC models in H2 are rather close to each other in terms of their BIC values. One of the top competitors to the model described above is a flexible DCC(1,1) model with a block structure of 1-2 kg and 2+ kg. The block structure is similar to that obtained in H1 in that it is the smallest fish that are the "outlier" among the weight classes. However, only the 1-2 kg weight class is separated from the remaining weight classes in H2. Meanwhile, the 2-3 kg weight class now belongs to the main block in terms of conditional correlations' dynamics. This change from H1 to H2 is in line with the development of the volume distribution across weight classes over time (see Figure 4); the 2-3 kg fish have become more abundant in the latter years in our sample.

5 Discussion

I have analyzed salmon spot price volatility over 1995-2013, comparing its magnitude and time patterns across the different weight classes. The first fundamental finding is a marked increase in volatility in all weight classes since around 2006, in agreement with Oglend (2013) observing a similar feature of the average spot price. This falls in line with several important structural changes in the salmon farming industry, such as the introduction of the maximum allowable biomass (MAB) restriction in 2005 and the opening of the salmon futures and options exchange "Fish Pool" in 2006. Having distinguished two sub-periods, from 1996 week 1 to 2005 week 45 (H1) and from 2007 week 27 to 2013 week 13 (H2), I proceed with modelling volatility in the two sub-periods separately.

The conditional mean and conditional variance dynamics of returns differs both qualitatively and quantitatively between the two subsamples. The autoregressive behavior of the conditional means and the ARCH(1) type of volatility clustering in H1 is replaced by moving average patterns in the conditional means and ARCH(3) effects in the conditional variances in H2. However, an autoregressive process often has an alternative moving-average representation, and vice versa, the difference between the two being rather technical. Examination of impulse responses in the conditional mean models reveals that the patterns in H1 are quite similar to those in H2, despite the apparent differences in model representations. (Impulse response graphs are available on request.) The ARMA patterns account for a larger proportion of the return variability in H1 than in H2 (except for fish of the extreme sizes). Combined with the increase in volatility, this suggests that the price behavior became more erratic after 2006. The high volatility and the low explanatory power of the conditional mean models in H2 imply that salmon market participants face extensive price uncertainty, which strengthens the demand for hedging and provides room for speculation using salmon derivatives at "Fish Pool".

The conditional correlations of shocks to returns are mostly positive throughout H1 and H2; however, the magnitude of the correlations increases markedly in H2 as compared with H1. Subsequently, the room for diversification targeted to reduce price uncertainty has shrunk. For example, compare the price uncertainty of a mixed portfolio of half small (say, 2-3 kg) and half large (say, 6-7 kg) fish against a uniform portfolio of medium-size fish. The mixed portfolio yields a smaller reduction in the price uncertainty over the uniform portfolio in H2 as compared with H1.

Consider fundamental reasons for the increase in correlations. The reduced supply from Chile during the ISA crisis increased the demand for Norwegian salmon across all the weight classes since 2007. The MAB restriction in 2005 may have changed the production schedule of different weight classes so that the weight distribution of available fish at any time in a year has changed, potentially leading to increased correlations. The introduction of the salmon futures and options exchange in 2006 increased the availability of timely information on both spot and futures prices. This information regards primarily the volume-weighted average price, thus it may tend to affect all weight classes in the same way, thereby increasing correlations. (Nevertheless, futures prices do carry weight-class-specific information given the known seasonality in volume distribution across the weight classes.)

Both the conditional mean and the conditional variance regularities differ also across weight classes, more so in H1 than in H2. For example, the conditional mean patterns are essentially the same for the few medium-size weight classes in H1 while the extreme sizes have their own distinct regularities. Small fish have larger ARCH effects than other weight classes in H2. The conditional correlations are larger between the neighboring weight classes than between distant weight classes. The latter is only natural as fish of similar size are closer substitutes than fish of dissimilar size.

There is a close correspondence between the dynamics of the three most voluminous weight classes throughout the sample period. Looking at both the univariate ARMA-GARCH models and the conditional correlation estimates for the 3-6 kg weight classes in H1 and especially H2, we could say that their return dynamics are essentially the same. Hence, the three weight classes could be treated as one without significant loss of precision. As noted by one of the referees, the relative similarity of the return regularities across weight classes may be providing support for the well-functioning of the salmon futures market.

The coincidence of the opening of “Fish Pool” with the increase in volatility, more erratic return behavior and higher correlations across weight classes is a curious one. All these changes are beneficial for the futures and options exchange, as discussed above. One can only wonder whether the launch of “Fish Pool” was a result of or perhaps a cause for these developments. Analysis of the information flows between the futures and spot markets and the role of both markets in the price discovery process could help answer this question, which is left for future studies.

My findings differ in two important respects from Oglend and Sikveland (2008) and Oglend (2013). First, I contrast the price dynamics in H1 to H2 and find numerous differences between the two. Second, I provide a finer picture of salmon price volatility by using weight-class-specific data. It elucidates the differences and similarities between the price behavior of small and large fish, abundant and scarce weight classes. Hence the findings offer relevant and updated information for production and risk management decisions in the salmon industry and trade.

6 Acknowledgements

I thank Ole Gjølberg, Mindaugas Bloznelis, Per Bjarte Solibakke, Remigijus Leipus, Virmantas Kvedaras, Timo Teräsvirta, members of the Norwegian Center for Commodity Market Analysis and participants at 3rd Energy Finance Christmas Workshop in Oslo for valuable discussions and feedback. Also, comments by anonymous referees have helped improve the paper considerably. Financial support of the Norwegian Research Council is gratefully acknowledged.

References

AARSET, B. & JAKOBSEN, S.-E. 2009. Political regulation and radical institutional change: the case of aquaculture in Norway. *Marine Policy,* 33**,** 280-287.

ANDERSEN, T. G. & BOLLERSLEV, T. 1998. Answering the skeptics: yes, standard volatility models do provide accurate forecasts. *International Economic Review,* 39**,** 885-905.

ASCHE, F. & GUTTORMSEN, A. G. 2001. Patterns in the relative price for different sizes of farmed fish. *Marine Resource Economics,* 16**,** 235-247.

ASCHE, F. & GUTTORMSEN, A. G. 2009. Reguleringer og produksjonspraksis. *Norsk Fiskeoppdrett.*

BAUWENS, L., LAURENT, S. & ROMBOUTS, J. V. 2006. Multivariate GARCH models: a survey. *Journal of Applied Econometrics,* 21**,** 79-109.

BERGFJORD, O. J. 2007. Is there a future for salmon futures? An analysis of the prospects of a potential futures market for salmon. *Aquaculture Economics & Management,* 11**,** 113-132.

BILLIO, M., CAPORIN, M. & GOBBO, M. 2006. Flexible dynamic conditional correlation multivariate garch models for asset allocation. *Applied Financial Economics Letters,* 2**,** 123-130.

BOLLERSLEV, T. 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics,* 31**,** 307-327.

BRÆKKAN, E. H. 2014. Disentangling supply and demand shifts: the impacts on world salmon price. *Applied Economics,* 46**,** 3942-3953.

BRÆKKAN, E. H. & THYHOLDT, S. B. 2014. The bumpy road of demand growth—an application to atlantic salmon. *Marine Resource Economics,* 29**,** 339-350.

BROWN, M. B. & FORSYTHE, A. B. 1974. Robust tests for the equality of variances. *Journal of the American Statistical Association,* 69**,** 364-367.

CAPORIN, M. & MCALEER, M. 2013. Ten things you should know about the dynamic conditional correlation representation. *Econometrics,* 1**,** 115-126.

CAPPIELLO, L., ENGLE, R. F. & SHEPPARD, K. 2006. Asymmetric dynamics in the correlations of global equity and bond returns. *Journal of Financial Econometrics,* 4**,** 537-572.

CONT, R. 2001. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance,* 1**,** 223-236.

DAHL, R. E. & OGLEND, A. 2014. Fish price volatility. *Marine Resource Economics,* 29**,** 305-322.

ENGLE, R. 2002. Dynamic conditional correlation: a simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics,* 20**,** 339-350.

ENGLE, R. F. 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica: Journal of the Econometric Society,* 50**,** 987-1007.

GHALANOS, A. 2013. rugarch: univariate GARCH models. 1.2-0. ed. Vienna, Austria: R Project for Statistical Computing.

INCLAN, C. & TIAO, G. C. 1994. Use of cumulative sums of squares for retrospective detection of changes of variance. *Journal of the American Statistical Association,* 89**,** 913-923.

JENSEN, B.-A. 2013. *No immediate relief in sight for high salmon prices* [Online]. Available: http://www.intrafish.com/news/article1373426.ece [Accessed 2014-02-28].

LARSEN, T. A. & ASCHE, F. 2011. Contracts in the salmon aquaculture industry: An analysis of Norwegian salmon exports. *Marine Resource Economics,* 26**,** 141-150.

LEVENE, H. 1960. Robust tests for equality of variances. *In:* OLKIN, I. (ed.) *Contributions to Probability and Statistics: Essays in Honor of Harold Hotelling.* Palo Alto, Calif.: Stanford University Press.

MARINE HARVEST 2013. Salmon farming industry handbook 2013. Bergen: Marine Harvest.

MARTÍNEZ-GARMENDIA, J. & ANDERSON, J. L. 1999. Hedging performance of shrimp futures contracts with multiple deliverable grades. *Journal of Futures Markets,* 19**,** 957-990.

OGLEND, A. 2013. Recent trends in salmon price volatility. *Aquaculture Economics & Management,* 17**,** 281-299.

OGLEND, A. & SIKVELAND, M. 2008. The behaviour of salmon price volatility. *Marine Resource Economics,* 23**,** 507-526.

SILVENNOINEN, A. & TERÄSVIRTA, T. 2009. Multivariate GARCH models. *In:* ANDERSEN, T. G., DAVIS, R. A., KREISS, J.-P. & MIKOSCH, T. V. (eds.) *Handbook of financial time series.* Berlin: Springer.

SOLIBAKKE, P. B. 2012. Scientific stochastic volatility models for the salmon forward market: forecasting (un-) conditional moments. *Aquaculture Economics & Management,* 16**,** 222-249.

STRAUME, H.-M. 2014. Currency invoicing in Norwegian salmon export. *Marine Resource Economics,* 29**,** 391-409.

WOOLDRIDGE, J. M. 1991. On the application of robust, regression-based diagnostics to models of conditional means and conditional variances. *Journal of Econometrics,* 47**,** 5-46.

Paper 2

Short-term salmon price forecasting

Daumantas Bloznelis[1]

School of Economics and Business

Norwegian University of Life Sciences

Abstract

Price forecasting has fundamental value in the highly volatile market of farmed Atlantic salmon. Accurate price predictions could reduce uncertainty and aid planning decisions of salmon farmers, processors and other market participants. I forecast salmon spot price by univariate and multivariate time series models, artificial neural networks and k-nearest neighbours method using weekly data for 2007-2014. A naïve seasonal forecast reduces the mean absolute error of a naïve no-change forecast by up to 5% and predicts the direction of the price change correctly up to 65% of the time for 1-5 weeks ahead. However, no method seems to consistently improve upon the naïve seasonal forecast, lending support to a hypothesis of weak form efficiency of the salmon spot market.

Keywords: salmon price, price forecasting, ARIMA, ARFIMA, ARARMA, VAR, VECM, artificial neural network, k-nearest neighbours, subset selection, regularization

[1] Mailing address: School of Economics and Business, Norwegian University of Life Sciences, P.O. Box 5003, NO-1432 Ås, Norway. E-mail: daumantas.bloznelis@nmbu.no.

1 Introduction

Forecasting the future spot price of salmon is a bread-and-butter activity of salmon market participants. Salmon farmers tailor their slaughtering schedule to match the profit-maximizing production path, which depends crucially on the expected trajectory of the spot price. Salmon processors plan their operations relying on expectations of their input cost level that is determined by the spot price of salmon. Salmon exporters act as intermediaries between farmers and processors or retailers, and their profits vary with the development of the spot price. Traders of salmon futures contracts make or lose money depending on how good a forecast of the future spot price they have. Price volatility measured by the standard deviation of logarithmic returns on the weekly spot price has increased substantially since around 2006 (Bloznelis, 2016; Oglend, 2013). It was greater than that of any other soft commodity over 2007-2013[2] (Bloznelis, 2016). The high volatility is perceived negatively by the market participants (Jensen, 2013). Being able to forecast the salmon price with reasonable accuracy would help reduce uncertainty and aid planning; it could be an alternative to price hedging by forward or futures contracts.

Salmon farming is a major industry in Norway with output slightly above one million tons valued at around NOK 45bn in 2014 (The Fish Site, 2015)[3]. The production is predicted to expand by 3% annually over 2013-2020 (Marine Harvest, 2014b). Norway is the largest farmed salmon producer in the world with over 60% of the global production volume as of 2014 (Marine Harvest, 2014a); it exports nearly all its production and is the major salmon supplier for Europe, the largest salmon consumer continent. Other large salmon farming countries are Chile, Canada and United Kingdom. Norwegian salmon is traded in the spot market and using forward contracts; the exact volume shares of the two trade modes are unknown but are of comparable magnitude.

Salmon spot price results from the interplay of supply and demand. The supply volume is limited by the standing biomass and slaughterhouse capacity from above, and natural, biological and regulatory factors from below. Biomass is the total amount of fish in the farming facilities. It evolves relatively slowly due to continued slaughtering throughout the year and natural limitations on the speed of growth of salmon in the seawater. The speed of growth depends on seawater temperature and feeding intensity. Given a fixed amount of feed, salmon gains weight more rapidly at favourable temperatures but may gain very little or no weight at all when the water is too cold or too warm. Biomass is highly seasonal with peaks reached around September each year (Marine Harvest, 2014a).

---

[2] Spot price was used for salmon while front month futures prices were used for other commodities as the most representative measure of the spot price.

[3] To obtain production volume from export volume, I add 3-4% which accounts for the local consumption (Marine Harvest, 2014a).

Industry regulations limit the maximum biomass at each farming site; once the limit is hit, the farmer must start slaughtering or otherwise pay a fine.

Among the biological factors, infectious diseases and parasites play an important role. When the fish cannot be cured or this is too expensive, they are slaughtered and sold (most of the salmon diseases are not dangerous for humans). For example, an outbreak of infectious salmon anemia (ISA) in Chile decimated the salmon biomass in 2007-2010. This caused considerable fluctuations in supply and changed the global trade flows of salmon; see e.g. Asche et al. (2009).

Sexual maturation is another determinant of salmon supply. It takes place once the fish reach a certain age and results in inferior flesh quality; hence, farmers aim to slaughter the fish before it occurs. Sexual maturation is a highly seasonal phenomenon normally happening in late summer and autumn, but it may be sped up or delayed depending on feeding and other factors (Pall et al., 2006).

Salmon demand fluctuations are mainly seasonal. Christmas and Easter are prime examples of increased demand periods. In the long run, consumer tastes and purchasing power are the main drivers of demand.

In addition to the fundamental factors, currency exchange rate is an important element in the supply-demand relationship in the spot market for salmon. Almost all the Norwegian farmed salmon is exported (Marine Harvest, 2014a); while consumers face retail prices in foreign currencies, producers’ home currency is Norwegian kroner. EUR/NOK exchange rate has been found to be an economically significant determinant of the spot price in 2007-2013 (Bloznelis, 2016; see also Straume, 2013). If the salmon price does not immediately and fully adjust to the fluctuations in the EUR/NOK rate, the latter may be useful in forecasting the spot price in the short run.

With the launch of the Fish Pool futures and options exchange for salmon in mid-2006, a new medium of price discovery appeared. Daily price quotes on monthly contracts from one through 60 months ahead became publicly available. The trade on Fish Pool generates new information on the future development of the salmon spot market. Since price quotes for each contract have to be published even on days with no trades, the analysts at Fish Pool use their best judgment to reflect any relevant information in making up the quotes. This helps in timely dissemination of information and may influence the price formation in the spot market.

Another relevant source of information is the share prices of the salmon farming companies. Many of the major Norwegian salmon farming companies are listed on the Oslo Stock Exchange. The share price information is freely available with only a negligible time lag. The spot price of salmon affects the revenues of the farming companies, and thus it should be reflected in the share prices. Thus the share prices may be a medium of salmon price discovery. Carreira (2013) finds supporting evidence using share prices of a few large salmon farming companies in 2007-2013.

For a salmon farmer, both short-run and long-run price forecasts are important. In the short term, e.g. days or weeks, slaughtering volume may be increased or decreased within bounds of the harvest-ready biomass and within the capacity of the slaughtering plants. Feeding intensity is perhaps the only means to alter this available volume to a limited extent. In the long run, i.e. beyond one year, pen stocking schedule and thus production volume can be manipulated to a large extent. Long-term price forecasts are provided by a few private companies such as Kontali Analyse and Capia, by Oslo-Stock-Exchange-listed salmon farming companies in their quarterly financial reports and by bank analysts following the salmon farming sector. Meanwhile, short-term forecasts do not seem to be available, aside from the ultra-short term forecasts in Friday afternoon surveys by the fish news portal IntraFish[4]. Seeing this gap, I will focus on the short-term price forecasts.

Previous studies of salmon price forecasting are scarce. Lin et al. (1989) considered Norwegian salmon prices in the U.S. and the European Community (E.C.) using monthly data for 1983-1987. A structural dynamic simultaneous equation model together with Box-Jenkins time series techniques was used to predict the monthly price of salmon in the U.S. and the E.C. for 1989 through 1992. The forecasts were not evaluated as they targeted future prices that were unavailable at the time of publishing the article. Vukina and Anderson (1994) forecasted wholesale prices of five wild salmon species in Japan employing state space modelling techniques on monthly data for 1978-1991. Although they obtained gains over naïve no-change forecasts, their forecast accuracy varied widely across species and forecast horizons. Also, their hold-out sample consisted only of six points which is too few for making reliable inference. Asche et al. (2005) showed that wild and farmed salmon prices are cointegrated and law of one price holds using data for Japan; hence, the study of Vukina and Anderson (1994) is relevant with regards to the farmed salmon price. Gu and Anderson (1995) used monthly data for 1986-1993 in a state space model, similar to that of Vukina and Anderson (1994), to forecast both wild and farmed salmon prices in the U.S. wholesale market. Predictions for three through twelve months ahead were satisfactory in terms of mean absolute percentage error and delivered around 60% to 67% correct forecasts of the direction of price change. Finally, Guttormsen (1999) forecasted weekly prices of farmed Atlantic salmon organised into weight classes using data for 1992-1997. ARIMA, VAR and Holt-Winters exponential smoothing methods were used, among other. Four-, eight- and twelve-week forecasts were all rather accurate in terms of mean absolute percentage error and proportion of accurate forecasts of the direction of price change. However, the author notes that even the best model did not improve upon a naïve forecast when the most representative weight classes were considered.

[4] Internet address www.IntraFish.com

All the above studies are outdated from the empirical point of view. Thus I find the time ripe for an update. I will explore a number of forecasting techniques in pursuit of the best salmon price prediction for one to five weeks ahead. I will assess forecast accuracy using absolute and relative measures and will test forecast optimality. Implications of findings with regards to the salmon market participants will be discussed.

The remainder of the paper is organized as follows. The next chapter motivates the selection of variables and presents the data. Chapter 3 provides descriptions of forecasting techniques and forecast evaluation methods. Chapter 4 explores the forecasting results and compares them across the techniques. A discussion on the implications of findings concludes.

2 Data

The following variables are used in this study: salmon spot price, salmon export volume, share prices of salmon farming companies on the Oslo Stock Exchange, EUR/NOK exchange rate and salmon futures prices. Data on the remaining factors that may influence the salmon price is either unobservable or not publicly available at a weekly or higher frequency. Therefore, their effects will be implicitly incorporated in the models' seasonal components and error terms.

2.1 Spot price

Weekly data on the average spot price of salmon is publicly available for 1995 and onwards[5]. I study the period from 2007 week 1 to 2014 week 39 (2007w1-2014w39); see Figure 1. I do not use earlier data so as to match the available time periods of other variables and also since a structural change seems to have occurred in the spot price series around 2006 (Bloznelis, 2016). A fundamental change also occurred between 2013w13 and 2013w14 when the so-called NOS price was replaced by the so-called NASDAQ price. I will now briefly review the differences between the two as described in Fishpool (2014).

NOS survey prices are exporters' buying prices, i.e. prices paid to salmon farmers. Only representative exporters are included in the survey. The exporter must purchase at least two truck-loads (around 40 metric tons) of salmon per week, on average. The exporter must sell more than half of the total volume as gutted salmon to companies outside its own group. Only between-company sales are reported, where companies do not belong to the same owner. Prices are to be agreed upon on the spot, as opposed to forward contracting. (In practice, the prices are normally agreed upon a few days ahead of delivery rather than on the delivery day; they are nevertheless called spot prices.) With regards to reporting, slaughtering day is the decisive characteristic: the reported transaction is

[5] The data can be accessed at https://salmonprice.nasdaqomxtrader.com/public/home?1-1.ILinkListener-loginMenu-downloadIndexHistoryLink and https://salmonprice.nasdaqomxtrader.com/historicalNOSprices.xls

Figure 1 Salmon spot price: original (left), seasonally adjusted (middle), seasonal component (right)

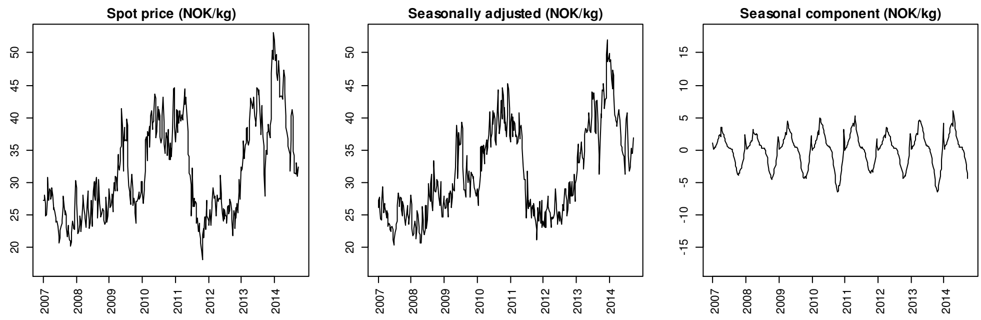

assigned to the week in which slaughtering took place. When published, NOS prices partly reflect the spot market prices as of the last week and partly the last week's deliveries' prices pre-agreed on Thursday or Friday the week before.

Meanwhile, NASDAQ survey prices are exporters' selling prices, i.e. prices received from buyers outside Norway. The NASDAQ survey inclusion conditions are different from those of the NOS survey. The exporters must sell at least five truck-loads (around 100 metric tons) of salmon per week, on average. Within-group as well as between-group sales may be reported. Only the export to Europe is included, while air-freight export to overseas markets is not. Customs invoice date (as opposed to slaughtering date) determines which week the transaction will be assigned to in the NASDAQ survey. There is a shorter lag between the actual transactions taking place and the prices being published; the NASDAQ salmon price is around four to five days old when reported. In sum, the NOS and the NASDAQ prices reflect a slightly different underlying object. Nevertheless, I will use them both (the NOS price for 2007w1-2013w13 and the NASDAQ price for 2013w14-2014w39) as measures of the spot price of salmon.

The difference in the treatment of dates, i.e. how a transaction is assigned to a calendar week, is important when forecasting the salmon price based on historical data. An observation from the NOS series points to a more distant past than one from the NASDAQ series, the difference being around five to six days. That is, if the NOS and the NASDAQ reporting were running simultaneously, the NASDAQ price for 2007w1 would likely match the NOS price for 2007w2 more closely than the NOS price for 2007w1. Therefore, given the last available NOS observation, forecasting one step ahead essentially means forecasting the past; one week ahead of around ten days ago is around three days ago. Meanwhile, forecasting the NASDAQ price one step ahead corresponds to forecasting the future; one week ahead of around four to five days ago is three to two days into the future.

If the NOS/NASDAQ weekly reports were the timeliest source of price information in the salmon spot market, the discrepancy between the two would likely yield different price discovery processes under the NOS reporting than under the NASDAQ reporting; this would need to be taken into account when modelling and forecasting the salmon price. However, it is inconceivable that the market participants do not get timelier price information, e.g. by trading and negotiating prices themselves. This is supported by the observation that the share prices of the salmon farming companies do not seem to react to the weekly NOS and NASDAQ reports. Meanwhile, they do seem to be moved by certain other reports containing timely information, such as the traders' survey published by the "IntraFish" and TDN every Friday afternoon.

The NOS and the NASDAQ prices were calculated simultaneously for 37 weeks in 2012-2013. The average difference between the NOS and NASDAQ prices was NOK 0.75/kg. I subtract NOK 0.75/kg from the NASDAQ price to be able to concatenate the two series. The same approach was used by the Fish Pool futures and options exchange for adjusting the Fish Pool Index at the time of transition from the NOS to the NASDAQ prices.

In addition to the average spot price, weight-class-specific spot prices are also available. Salmon is divided by fish size into seven weight classes, 1-2 kg, 2-3 kg, 3-4 kg, 4-5 kg, 5-6 kg, 6-7 kg and 7+ kg fish. The weight-class-specific prices lend themselves easily to operations' management by salmon farmers, processors, and other market participants. However, the distinction between the weight classes is of limited importance in terms of forecasting; the short-term movements in the weight-class-specific prices are highly correlated with the movements in the average price (Bloznelis, 2016). The correlations between seasonally-adjusted logarithmic returns vary between 0.93 and 0.99 for all the weight classes except for the scarcest class of the small 1-2 kg fish; the correlation for this weight class is still rather high at 0.77 (ibid.). Also, modelling the prices of all the seven weight classes instead of just one translates to increasing the number of dependent variables from one to seven. This would increase the model complexity and would aggravate the curse of dimensionality. Therefore, I leave it for future studies. Ultimately, if a need arises to forecast the price of a specific weight class, it could be done by modelling the spread between the weight class of interest and the average price, and adding it to the forecast of the average price.

Seasonality is a key characteristic of salmon production and is due to the seasonality in the supply and demand factors discussed in Chapter 1. The seasonality of supply does not match the seasonality in demand, which produces seasonal patterns in price. Modelling the seasonality for weekly time series is known to be tricky. The seasonal period is long and there is a fractional number of weeks in a year which limits the choice of available techniques and excludes the most popular ones such as dummy variables, seasonal ARIMA model and STL decomposition by Cleveland et al. (1990). Here I follow an approach proposed by Hyndman (2014) of using a regression with ARMA errors, with Fourier terms as regressors. The number of Fourier terms could be up to 26 pairs for weekly data; however, a parsimonious model may exclude the higher order terms (the ones generating the high-frequency oscillations) so as to trade off an increase in the model misspecification bias against a decrease in estimation imprecision. This is formalized by choosing between none through 26 pairs of Fourier terms using the corrected Akaike's information criterion (AICc). The spot price and the export volume exhibit sharp, large movements around Christmas and Easter. I therefore also add four dummy variables to indicate weeks right before and after Christmas and another four for Easter. Note that even though Christmas has a fixed date in the year, that date may belong to different weeks over different years if the weeks are counted from Monday to Sunday. The Fourier terms together with the Christmas and Easter dummies constitute the seasonal pattern and will be subtracted from the original series, while the remainder will be used for further modelling. The seasonal component and the seasonally-adjusted salmon spot price for the whole period 2007w1-2014w39 are depicted in Figure 1.

2.2 SSB export volumes

The second variable of interest is salmon production volume. While production data is not available at a weekly frequency, export volume can be a good proxy for it as the local consumption only makes up some 3-4% of the total production (Marine Harvest, 2014a). The total export volume of fresh farmed Atlantic salmon provided by Statistics Norway (SSB)[6] represents all farmed Atlantic salmon export out of Norway. The data is obtained from the Norwegian Customs' electronic information system for the exchange of customs declarations between businesses and the Norwegian Customs (TVINN). The volume for the last whole week is published every Wednesday; hence, the report covers the period of nine to three days before the publication date. Both fresh and frozen salmon volumes are available, but the latter are relatively small and only the former are used in this study. The data exhibits a time trend and strong seasonality. Seasonal adjustment is performed in the same way as in the case of the spot price. The original export volume series, the seasonally-adjusted series and the seasonal component are depicted in Figure 2.

2.3 Salmon futures

Daily price quotes of salmon futures contracts with maturities of one through 60 months are available at the Fish Pool website[7] for mid-2006 and onwards. I transform the data to a weekly frequency by taking the last daily observation of each week. Although taking a weekly average would be more directly comparable to the spot price and the export volume data, the last observation of the week provides the most up-to-date information and thus is more relevant for forecasting. I then seasonally adjust the weekly futures prices using a similar approach as with the spot price. Necessary changes are made due to the seasonal pattern being monthly rather than weekly as the underlying of

[6] The data can be accessed at http://www.ssb.no/en/utenriksokonomi/statistikker/laks?fane=om
[7] The data can be accessed at http://fishpool.eu/fishpool/servlets/newt/public/excel?type=historic

Figure 2 Salmon export volume: original (left), seasonally adjusted (middle), seasonal component (right)

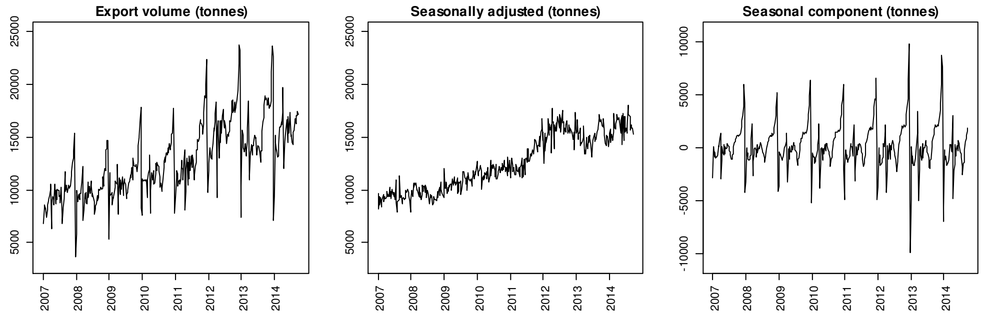

a futures contract is the average salmon price over a month. Figure 3 shows the original series together with its seasonally-adjusted version and the seasonal component.

The futures contracts exhibit large price changes every time the underlying of the contract changes; e.g. the underlying of the front-month contract turns from the average salmon price over January, 2014, to the average salmon price over February, 2014. This is the well-known "rollover effect", an artefact due to the construction of the data series. However, it is a nuisance; while normally the price changes reflect the evolving price expectations of a fixed underlying, the price change on a rollover date comprises a combination of a change of the underlying itself and a change in the price expectation of the fixed new underlying (February, 2014, in the above example). Rollover effects need to be adjusted for. First, I transform the futures prices to the corresponding logarithmic returns (log-returns). Then I replace the observations on rollover dates by the log-returns for the new underlying months. After that I transform the data from log-returns back to levels. Note that this rollover-adjusted variable does not have a clear economic interpretation. Still, it may be useful for forecasting.

Even after the adjustment for the rollover effects, the futures series are not directly relevant for forecasting the spot price in the short run. The changes in the price expectations for a given future month may be of limited relevance to the price expectations of the next few weeks. I suggest that a more relevant transformation of the futures prices can be constructed. The idea is to match the underlying of the futures contract with the underlying of the spot price to be forecasted. Suppose that someday in 2014w4, e.g. Monday, 27th January 2014 we want to forecast the spot price of the next week (2014w5) that corresponds to 3rd through 9th February 2014. Obviously, the price of the February futures contract (which covers 2014w5-2014w8) is more relevant than the price of the January futures contract (2014w1-2014w4). Thus the next month contract is preferred to the front month contract as an explanatory variable in a forecasting model. Meanwhile, this would not be true

Figure 3 Front month futures price: original (left), seasonally adjusted (middle), seasonal component (right)

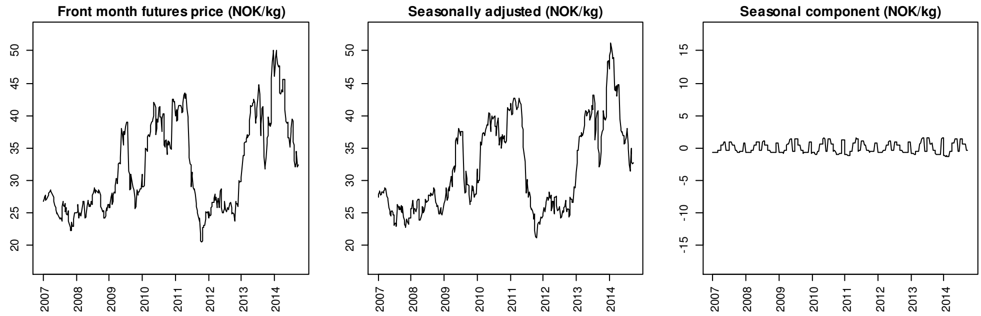

if we stood one week before in calendar time; on Monday, 20th January 2014 trying to forecast spot price of 2014w4 (27th January through 2nd February 2014), the front month futures contract (2014w1-2014w4) is more relevant than the next month futures contract (2014w5-2014w8). Thus the choice of the explanatory variable depends on where we are in the calendar.

Following this logic, I construct an artificial variable which is a combination of the front month and the next month futures contracts such that it would exactly fit the need of forecasting one week ahead. It makes sense to consider log-returns instead of log-levels here. The new variable equals the log-return on the front month contract everywhere except for the last week of each trading month (which is similar but not exactly the same as calendar month); there it equals the log-return on the next month's contract. When transformed from log-returns back to levels, I call this variable the futures predictive variable. It is graphed in Figure 4 together with its seasonally-adjusted counterpart and the seasonal component.

2.4 Share prices

The five largest salmon farmers in Norway are Marine Harvest (Oslo Stock Exchange ticker MHG), Lerøy Seafood Group (LSG), SalMar (SALM), Cermaq (CEQ) and Grieg Seafood (GSF) (Marine Harvest, 2014a). I exclude Cermaq due to salmon farming in Norway being only a small fraction of its business for the most part of the sample period. Meanwhile, the other four companies have salmon farming in Norway as their main activity. As of 2013, together they accounted for around 53% of the salmon production in Norway (Marine Harvest, 2014a), although this share used to be lower in earlier years.

Daily closing prices of the four shares are made publicly available by "Netfonds"[8]. I have manually adjusted the series for dividend payments and a reverse split of the MHG share in 2014. For

[8] Internet address www.Netfonds.no

Figure 4 Futures predictive variable: original (left), seasonally adjusted (middle), seasonal component (right)

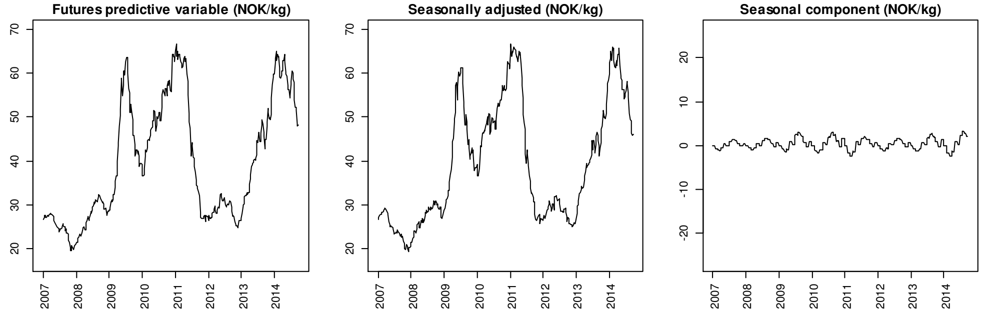

Figure 5 Salmon farming companies' share prices

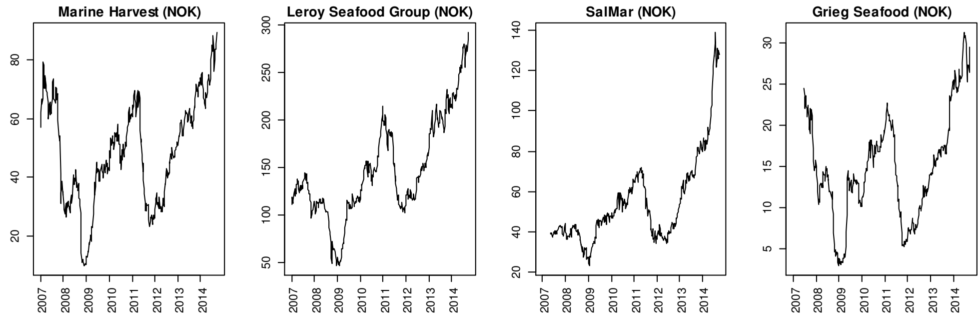

compatibility with the spot price data, I transform these to a weekly frequency by picking the last observation of each week. The transformed prices are plotted in Figure 5.

I construct an index corresponding to an equally-weighted portfolio of the four shares. Since SALM and GSF shares were listed on the Oslo Stock Exchange as late as May and June, 2007, respectively, I use only MHG and LSG shares for the first few months to construct the index. I include SALM into the index from the 21st day after the initial public offering (IPO), and I do the same with GSF. The first 20 days are not included to avoid any possible effects of the recent IPO. The index is normalized so that its logarithm equals one for the last week before 2007w1. In contrast to the individual share prices, the index should be less sensitive to company-specific issues and better reflect sector-wide factors such as the development of the salmon spot price. The index is depicted in Figure 6.

Figure 6 Salmon farming companies' share price index

The Oslo Seafood Index could be considered as an alternative to the index proposed above. The Oslo Seafood Index is a stock market index consisting of eleven companies from the Norwegian seafood industry (Oslo Børs). Its main drawback is the inclusion of companies such as Havfisk (a whitefish fishing company) or Hofseth Biocare (a fish oil and pharmaceuticals producer) that do not depend significantly on the spot price of salmon. Also, the historical price series for the index are not available before June 2010, i.e. it is three and a half years shorter than the salmon price series. Therefore, I use the Salmon farming companies' share price index described above.

2.5 Currency exchange rate

The last variable included in this study is the EUR/NOK exchange rate. The data is obtained from Oanda[9] and used as is (in logarithmic form). Figure 7 depicts the EUR/NOK rate development over 2007w1-2014w39.

Figure 7 EUR/NOK currency exchange rate

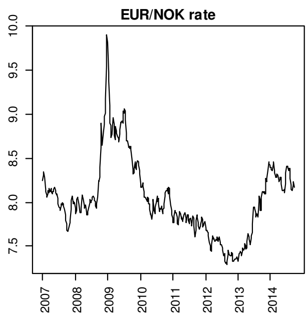

[9] The data can be accessed at http://www.oanda.com/currency/historical-rates/

Due to lags in reporting and the nature of weekly data, the whole set of the newest available spot price, export volume, futures price, share prices and EUR/NOK rate data may be outdated by roughly one to two and a half weeks, depending on the weekday. If forecasts are to be made using only the publicly available data, one-week-ahead forecasts will actually be forecasts of a past period. Two-week-ahead forecasts will be now-casts of the current complete week. Only the three-week-ahead forecasts will really be about the future, i.e. the next complete week. It is reasonable to expect the market participants to have more timely information on the salmon spot price and the export volumes, where the reporting lag is the largest. If a good estimate of the yet-unreported spot price and export volume data can be obtained, the one-week-ahead forecasts may become now-casts, while the two-weeks-ahead forecasts would already target the future. This should be kept in mind when assessing the forecast accuracy and considering the practical applicability of the forecasting methods.

2.6 Descriptive statistics

Graphs of the spot price, the export volume, the futures predictive variable, the share price index and the EUR/NOK rate are provided in Figures 1-7; descriptive statistics are provided in Table 1. All the data series are non-normally distributed as evidenced by the Shapiro-Francia, D'Agostino and Anscombe-Glynn tests (Shapiro and Francia, 1972; D'Agostino, 1970; Anscombe and Glynn, 1983).

Modelling will be done on the seasonally-adjusted logarithms or logarithmic returns of the data so as to bring skewness and excess kurtosis closer to zero, linearize any exponential growth over time and turn the potentially multiplicative seasonality and shocks into additive ones. Even though the skewness and the excess kurtosis of the logarithmically transformed data are closer to those of a normal random variable, logarithmic series are still non-normally distributed.

Analysis of stationarity or integratedness of the series will be useful when working with time series forecasting models. Results of the augmented Dickey-Fuller test (Said and Dickey, 1984) do not allow rejecting the null hypothesis that the variables are integrated (regardless of inclusion of a drift term in the test equations) and that their first differences are stationary. The more powerful modified Dickey-Fuller (DF-GLS) test (Elliott et al., 1996) confirms this finding (regardless of the type of GLS detrending being "constant" or "trend") for all variables but the export volume. Whether the export volume can be treated as a random walk is debatable. On the one hand, it is not a cumulative sum of past shocks because each generation of fish spends only a limited time in the seawater and takes away the effects of shocks from the volume series when harvested. However, even though natural shocks do not have an impact on fish that are yet to be released into the seawater, they might induce changes in the release schedule. Also, the export volume exhibits a slowly-changing (as opposed to linear) time trend which could be treated as a random walk for convenience. An

Table 1 Descriptive statistics of spot price, export volume, front month futures price, share price index and EUR/NOK rate

Weight class	Mean	Median	Min.	Max.	St. dev.	Coef. of variation	Skewness (D'Agostino)	Excess kurtosis (Anscombe)	Normality (Shapiro-Francia)
Spot price	31.68	29.09	18.17	53.17	7.39	0.23	0.55***	-0.70***	0.9414***
Export volume	12 756	12 302	3 704	23 674	3 392	0.27	0.47***	-0.03	0.9717***
Front month futures price	32.01	29.00	20.55	50.00	7.05	0.22	0.56***	-0.87***	0.9179***
Share price index†	3.03	2.92	0.86	6.37	1.14	0.38	0.70***	0.55**	0.9595***
EUR/NOK rate	8.04	7.98	7.30	9.90	0.44	0.05	0.93***	1.45***	0.9475***

Note: ***, ** and * mark statistical significance at 1%, 5% and 10% level, respectively. Statistical significance of skewness, kurtosis and normality is assessed by d'Agostino, Anscombe and Shapiro-Francia tests, respectively. † scaling of the share price index is arbitrary, as discussed in Chapter 2, Section 2.4.

ARIMA(3,1,0) model with a drift term seems to fit the export volume rather well. Thus I will treat the export volume series as a random walk with drift for simplicity. Another characteristic relevant for a system of integrated variables is cointegration; it will be discussed in the context of the VEC model in the next chapter.

Analysis of cross-correlations helps elicit lead and lag relationships between variables, which will be useful when considering multivariate forecasting models. Since the logarithmic variables are integrated, the cross-correlations are high but little informative. More revealing are the cross-correlations of their first differences (log-returns, or simply returns) as depicted in Figure 8.

In Figure 8, statistically significant cross correlations at positive lags indicate that the second variable in the title of the graph leads the first variable by a number of weeks equal to the lag number. Conversely, cross correlations at negative lags appear when the first variable in the title of the graph leads the second variable. Thus the significant cross correlations at the positive lags in row one of Figure 8 are likely to help predict the spot price. A few interesting observations can be distinguished. First, the spot price returns are autocorrelated at lag two. Second, the returns on the export volume of the most recent two to three weeks are significantly correlated with the spot price returns; contemporaneous correlation is also present but is of limited use for forecasting. Third, the share price index returns lead the spot price returns by one week, and the correlation is quite strong.

Figure 8 Cross correlations between weekly changes in seasonally adjusted variables

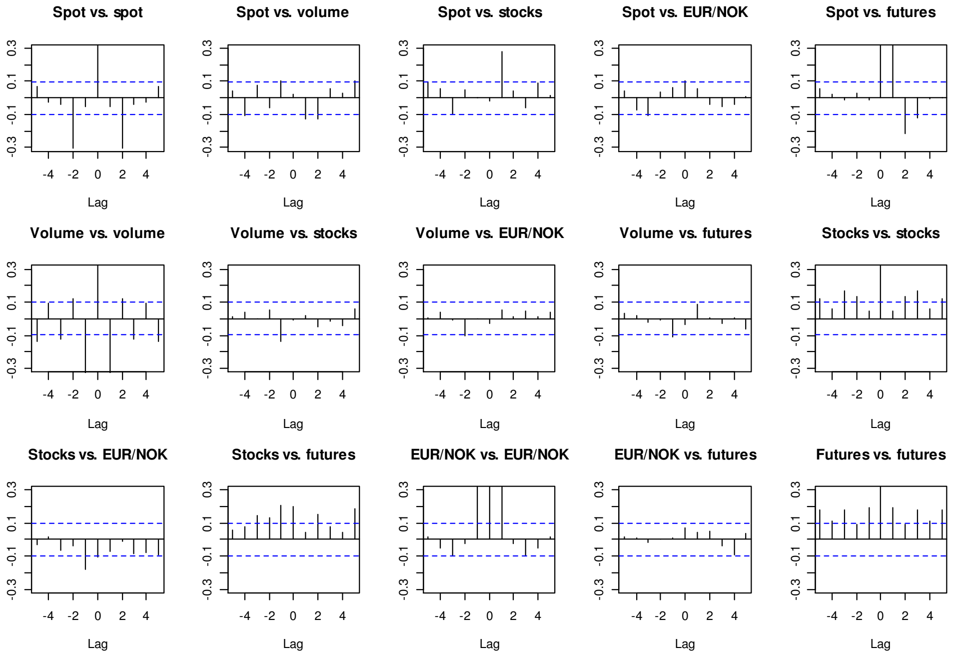

Note: cross correlations at positive lags indicate that the second variable in the title of the graph leads the first variable. E.g. the middle graph in the top row shows that the share price index (stocks) leads the salmon spot price (spot) by one week. Cross correlations at negative lags indicate that the first variable in the title of the graph leads the second variable. E.g. the second graph in the middle row shows that the export volume (volume) leads the share price index (stocks).

Fourth, the EUR/NOK rate returns are contemporaneously correlated with the spot price returns, but this is again of limited use for forecasting. Fifth, the returns on the futures predictive variable lead the spot price returns by one, two and three weeks. The significant lead-lag relationships in rows two and three of Figure 8 may also be used in multivariate models for forecasting more than one period ahead.

Note that the analysis of cross-correlations is based on the whole 2007w1-2014w39 period and therefore could not have been used for forecasting any period within the sample in real time. Therefore, these findings will not actually be used in modelling; the full sample analysis is provided here for completeness and as guidance for the future research. However, using a shorter sample ending on 2012w52 gives similar results which then are used for forecasting the periods starting with 2013w1 and beyond.

Also note that the cross correlations where the spot price returns are involved should be interpreted carefully for the period before 2013w14 due to the dating convention used by NOS. For example, if the share price index returns are found to lead the spot price returns by one week, this does not necessarily mean the correlation could be utilized to predict the spot price returns in real time. The apparent lead by one week may be due to the fact that the spot price recorded for a given week is partly determined the week before. Thus only two or more weeks' leads may be truly useful in the real time forecasting.

3 Methods

I will use a number of forecasting methods to predict the salmon spot price one to five weeks ahead. To realistically replicate the situation faced by a forecaster at a given point in time, I will split the sample into a period of available "past" data and a period of unavailable "future" data, or an "in-sample" and a "hold-out sample". The "in-sample" will be used to estimate candidate models and select the optimal one of them. Then the selected model will be employed to forecast the spot price h periods ahead, for $h = 1, \dots, 5$. This exercise will be repeated taking weeks 2013w1 through 2014w39 as the splitting points between the "in-sample" and the "hold-out sample". Forecasting performance will be assessed using the "hold-out samples". Note that the candidate model specifications are fixed over time, but the optimal model specification is reassessed with every new time period; such an approach is more computationally intensive than using a fixed model and only re-estimating the model's parameters. This is a realistic setting in that only the information available to the forecaster at the point of forecasting is used, with no adjustments for future information.

I start with six years (313 weekly observations) of data from 2007w1 through 2012w52 as the initial "in-sample" and produce forecasts as of 2012w52. Then I move the six-year window one step ahead by deleting the data for 2007w1 and appending a new data point of 2013w1. There are 92 windows from 2007w1-2012w52 through 2012w40-2014w39. The deletion of the first data point as a new data point is appended amounts to using a rolling window rather than an expanding window. It is justified when the characteristics of the data change rapidly enough over time. The trade-off is between a loss in estimation precision (due to a smaller sample size when deleting the first observation in the sample) and an improvement in the model's relevance (due to the earliest observation being too little representative of the prevalent relationships between the variables within the data window). The choice of the rolling window over the expanding window in this study is arbitrary. Conveniently, it yields lower computational time due to the non-increasing sample size. However, both options could in principle be considered. Alternatively, the choice could be based on cross validation (discussed below).

Conventional time series models produce one-step-ahead forecasts which can be appended to the data to iteratively produce h-step-ahead forecasts for $h > 1$. On the other hand, direct multi-step-ahead forecasting is also possible when regressors lagged by at least h periods are used.

Empirical research does not seem to favour direct multi-step forecasts over iterative ones or vice versa, see e.g. Marcellino et al. (2006) and Pesaran et al. (2011), and references therein; I choose the conventional iterative forecasts for simplicity. For nonparametric methods such as the k-nearest neighbours method and artificial neural networks, I use direct h-step-ahead forecasts which are perhaps more natural, although iterative one-step-ahead forecasts are also feasible.

3.1 Model selection via information criteria and cross validation

In a time series setting, selection of one model from a pool of candidate models is usually done in either of the following two ways. First, one may pick the model with the lowest AIC value. Second, one may use time series cross validation and select the model with the lowest root mean squared forecast error (RMSE). Asymptotically, minimizing the AIC is equivalent to minimizing the RMSE of the one-step-ahead out-of-sample forecasts (Konishi and Kitagawa, 2008, p. 249-250). (Compare to the cross-sectional setting where minimizing the AIC is asymptotically equivalent to minimizing the RMSE in leave-one-out cross validation (Stone, 1977)). Thus in large samples AIC-based selection is optimal with respect to the forecast RMSE.

Cross validation is a technique that uses part of the data (the "training sample") to estimate the candidate models and the remainder of the data (the "validation sample") to assess their performance, based on which the best model is selected; see e.g. Hastie et al. (2009, p. 222). The sample is usually split randomly multiple times into different training and validation samples to grant robustness to the validation results. In the time series setting, sample splitting cannot be done as flexibly as in the cross-sectional setting since the order of the observations matters. Hyndman and Athanasopoulos (2014) describe a cross validation technique for time series. They suggest using a rolling window inside the original sample. The entire rolling window except for the last observation serves as the training sample, while the last observation constitutes the validation sample. Training, validation and model selection is subsequently done in the same way as in the regular cross validation. If model selection is to be followed by testing its performance on new data, Hastie et al. (2009, p. 222) recommend splitting the original sample into 50% for training, 25% for validation and 25% for testing. Correspondingly, I use a four-year rolling window of weekly data for training and validation within a six-year sample. The latter is itself a rolling window within an eight-year original sample, where rolling is used for forecast evaluation ("testing").

Should the AIC or cross validation be preferred? The strengths of the AIC are, first, that it is simple to calculate and, second, that it uses the whole "in sample" rather than just the training part of it, and thus avoids bias towards more parsimonious models. The weakness is that over-fitting can be an issue especially across groups of models that have the same number of parameters; there the penalty for model complexity does not differ across the models, and minimizing the AIC amounts to maximizing the model likelihood. Also, application of the AIC to nonparametric methods such as the k-nearest neighbours or artificial neural networks is neither straightforward nor intuitive. Meanwhile, the strength of cross validation is that it selects the model by a natural criterion of minimizing the RMSE of the one-step-ahead forecasts. In small or medium-size samples, this is in contrast to the AIC-based selection as the asymptotic equivalence between minimizing the AIC and minimizing the forecast RMSE cannot be invoked. The weakness is that time series cross validation systematically favours too-parsimonious models. Given a pool of candidate model specifications, the optimal forecasting model generally depends on the sample size[10]. Cross validation uses training samples that are smaller than the whole "in-sample"; hence, richer models are at a disadvantage and are selected less frequently than would be optimal given the actual "in-sample". Cross validation is also subject to over-fitting when the number of candidate models is large relative to the size of the training sample. Following the arguments above, I choose to employ AIC-based selection for the time series models but cross validation for the k-nearest neighbours method and artificial neural networks. Cross validation is also used together with regularization for multiple time series models.

When logarithms instead of levels are employed in the forecasting models, the forecasts have to be transformed back to levels. Curiously, undoing the logarithmic transformation by exponentiation need not be an optimal solution. For example, the exponent of the mean of a normally distributed random variable is not equal to the mean of the exponent of that variable. Thus when the conditional mean is to be forecasted, exponentiation is generally suboptimal. However, Bårdsen and Lütkepohl (2011) show that it works fine in practical applications where model specification and estimation uncertainty play a role. When the logarithmic transformation of the data is non-normal, a smearing estimate could be used (Duan, 1983). The latter approach is rather tedious, thus I proceed with simple exponentiation instead.

3.2 Individual models

3.2.1 Univariate ARIMA

ARIMA stands for autoregressive integrated moving average. We say that a time series $\{y_t\}_{t=1}^{\infty}$ follows an ARIMA(p, d, q) process if

$$\Delta^d(y_t - \mu) = \varphi_1 \Delta^d(y_{t-1} - \mu) + \cdots + \varphi_p \Delta^d\left(y_{t-p} - \mu\right) + \varepsilon_t + \theta_1 \varepsilon_{t-1} + \cdots + \theta_q \varepsilon_{t-q}$$

where μ is the mean of the process and Δ^d represents differencing of order d. Differencing of order $d = 1$ is implicitly defined as $\Delta y_t := y_t - y_{t-1}$, while differencing of order $d > 1$ is defined

[10] I am not aware of any source that has explicitly mentioned this problem. Thus it may be a new contribution to the time series model selection and forecasting methodologies.

recursively as $\Delta^d y_t := \Delta^{d-1}(\Delta y_t)$. The simplest representative ARIMA model is ARIMA(1,0,1) also known as ARMA(1,1), with zero mean:

$$y_t = \varphi_1 y_{t-1} + \varepsilon_t + \theta_1 \varepsilon_{t-1}$$

An integrated ARIMA($p, 1, q$) process is obtained if $\{y_t\}_{t=1}^{\infty}$ are cumulative sums of all values of x_τ for $\tau = 1, \dots, t$ where $\{x_t\}_{t=1}^{\infty}$ is itself an ARMA(p, q) process. Such a process is integrated of order $d = 1$, which is the most common in practice. Higher integration orders $d > 1$ can be obtained by cumulatively summing $\{x_\tau\}_{t=1}^{\infty}$ where $\{x_\tau\}_{t=1}^{\infty}$ is an ARIMA($p, d-1, q$) process. Integratedness can be tested using unit root tests such as the ones mentioned in Chapter 2. Integrated ARIMA(p, d, q) processes are differenced d times to achieve stationarity and then modelled as ARMA(p, q) processes.

I determine lag orders p and q using the AIC and full subset selection. For a given sample, I estimate models of all possible lag order specifications up to a pre-specified maximum lag order (p_{max}, q_{max}) and select the model that minimizes the AIC. I do not anticipate distant lags to have substantial predictive power, thus I cap the maximum lag order at four for both the autoregressive and the moving-average part. This amounts to $2^{4+4} = 256$ subset models to choose from.

3.2.2 Univariate ARFIMA

Univariate autoregressive fractionally-integrated moving-average (ARFIMA) model pioneered by Granger and Joyeux (1980) generalizes the ARIMA model with respect to the order of integration. While the ARIMA model allows the process to be stationary or integrated of an integer order, the ARFIMA model fills the gap by allowing for a fractional order of integration. The general model formula is the same as that for ARIMA, but the differencing operator is extended to cover fractional orders of integration. A fractional difference d of a time series variable y_t is defined as follows:

$$\Delta^d y_t := y_t - d y_{t-1} + \frac{d(d-1)}{2!} y_{t-2} - \frac{d(d-1)(d-2)}{3!} y_{t-3} + \cdots$$

Fractional integration allows for presence of the so-called long memory. In an ARIMA($p, 0, q$) model the impact of innovations ε_t onto y_{t+h} for $h > 0$ decays exponentially in h, which is known as short memory. In ARIMA($p, 1, q$) model it does not decay at all, which is called infinite memory. Meanwhile, it decays slower than exponentially in an ARFIMA(p, d, q) model, which mimics the behaviour of some macroeconomic and financial variables such as inflation (Baillie et al., 1996) (Doornik and Ooms, 2004) or spot price of electricity (Koopman et al., 2007). I estimate the order of integration d using the “autoarfima” function in the “rugarch” package in R (Ghalanos, 2014).

I choose the autoregressive and the moving average lag orders just as in the case of the ARIMA model, and cap the maximum lag orders at four. The computational burden is increased over the conventional ARIMA model due to the estimation of the order of integration parameter d.

3.2.3 Univariate ARARMA

Autoregressive-autoregressive moving average (ARARMA) model proposed by Parzen (1982) is another generalization of the ARIMA model. The idea is to replace the differencing of an integrated time series by fitting an AR(τ) model with only the τ^{th} coefficient not equal to zero and keeping the residuals. The optimal lag τ is selected to minimize the loss

$$Err(\varphi(\tau),\tau) = \frac{\sum_{t=\tau+1}^{n}(y_t - \varphi(\tau)y_{t-\tau})^2}{\sum_{t=\tau+1}^{n} y_t^2}.$$

This is called detrending, or memory shortening. (The special case when $\tau = \varphi(\tau) = 1$ corresponds to the regular differencing as in the ARIMA model described above.) The residuals from the AR(τ) model are further modelled as an ARMA process. The guidelines for the ARMA model specification are outlined in Parzen (1982), but for simplicity the same specification method as in the case of ARIMA will be used. When forecasting with an estimated ARARMA model, I first make a forecast for the de-trended data, and then back up the original data by re-trending. The ARARMA method fared well in a forecasting tournament "M-competition" (De Gooijer and Hyndman, 2006; Makridakis et al., 1982; Meade and Smith, 1985).

Interestingly, the optimal lag order is found to be $\tau = 1$ for the salmon spot price in all the rolling windows, and the corresponding coefficient values φ equal unity up to three significant digits; other values of τ produce significantly higher $Err(\varphi(\tau),\tau)$. Thus the first step of ARARMA modelling essentially equals first differencing, and therefore ARARMA coincides with ARIMA in this particular case. Therefore, I do not assess ARARMA forecasts separately but say that they are essentially no different from those of the ARIMA model.

3.2.4 Vector ARIMA

A multivariate generalization of the ARIMA model is vector ARIMA, or VARIMA model. The formulation and interpretation follows straightforwardly from its univariate counterpart. A general form of the VARIMA model is

$$\Delta^d(\boldsymbol{y}_t - \boldsymbol{\mu}) = \Phi_1\Delta^d(\boldsymbol{y}_{t-1} - \boldsymbol{\mu}) + \cdots + \Phi_p\Delta^d(\boldsymbol{y}_{t-p} - \boldsymbol{\mu}) + \boldsymbol{\varepsilon}_t + \Theta_1\boldsymbol{\varepsilon}_{t-1} + \cdots + \Theta_q\boldsymbol{\varepsilon}_{t-q}$$

where $\boldsymbol{y}$ is a $K \times 1$ vector of the dependent variables, $\boldsymbol{\mu}$ is the mean of the process, Φ_1 through Φ_p and Θ_1 through Θ_q are quadratic $K \times K$ coefficient matrices, and Δ^d is order d element-wise differentiation operator. The simplest representative example is a VARIMA(1,0,1), or VARMA(1,1) for a bivariate time series $\begin{Bmatrix} y_{1,t} \\ y_{2,t} \end{Bmatrix}_{t=1}^{\infty}$ with zero mean:

$$\begin{pmatrix} y_{1,t} \\ y_{2,t} \end{pmatrix} = \begin{pmatrix} \varphi_{11} & \varphi_{12} \\ \varphi_{21} & \varphi_{22} \end{pmatrix} \cdot \begin{pmatrix} y_{1,t-1} \\ y_{2,t-1} \end{pmatrix} + \begin{pmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{pmatrix} + \begin{pmatrix} \theta_{11} & \theta_{12} \\ \theta_{21} & \theta_{22} \end{pmatrix} \cdot \begin{pmatrix} \varepsilon_{1,t-1} \\ \varepsilon_{2,t-1} \end{pmatrix}.$$

Unlike a set of univariate ARIMA models for each component of $\boldsymbol{y}$, VARIMA model allows for cross-equation relationships, such that the first component $y_{1,t}$ may be a function of a lagged value of the second component $y_{2,t-i}$ and a lagged value of innovation to a third component $\varepsilon_{3,t-j}$, for example. However, such a flexible model form comes bundled with estimation complexity. It is not uncommon that the estimation of VARIMA models suffers from convergence to local, rather than global, optimum and from surface of the likelihood being rather flat. Also, existence of equivalent representations of the same model becomes a relatively large problem; there are a number of possible representations, some more parsimonious than other, and the most parsimonious representation is cumbersome to derive; see Gilbert (1993) for an overview. Due to the difficulties listed above I will use a special case of the VARIMA model that is considerably easier to work with and to estimate.

3.2.5 VAR

A VAR model is a special case of VARIMA in which lagged innovations are not considered:

$$\Delta^d(\boldsymbol{y}_t - \boldsymbol{\mu}) = \Phi_1 \Delta^d(\boldsymbol{y}_{t-1} - \boldsymbol{\mu}) + \cdots + \Phi_p \Delta^d(\boldsymbol{y}_{t-p} - \boldsymbol{\mu}) + \boldsymbol{\varepsilon}_t.$$

Even if the true data generating process were of VARIMA type, it could be approximated arbitrarily well by a VAR model with sufficiently many lags, under certain regularity conditions; see Lütkepohl (2007, p. 531-553) for details and practical considerations. This would increase the number of parameters relative to the most parsimonious VARIMA representation but would also bring a considerable advantage of simple and fast estimation. A VAR model can be estimated consistently equation-by-equation by the OLS method. For a full unconstrained VAR model (where the lag order is the same in all equations and no linear or nonlinear parameter restrictions are imposed) this is also the efficient estimator; for a subset VAR model (where the lag order is not the same in all equations or some parameter constraints are imposed) OLS is still consistent but not efficient; the efficient estimator is the feasible generalized least squares (FGLS) estimator, which is barely more complex and still rather fast to compute. The simplicity of the VAR model estimation is in stark contrast to the cumbersome and computationally intensive fitting of the VARIMA model.

Lag order selection is a key step in VAR modelling. The number of possible different combinations of lag orders is vastly greater than in the case of the ARIMA model. The number of subset models grows very fast with the lag order p as well as with the number of the dependent variables K; more exactly, the number of subset models is $2^{K \cdot K \cdot p}$. For example, it exceeds $4 \cdot 10^{37}$ for a VAR system of five response variables and five lags in each equation. This makes estimation of all the available models computationally infeasible. Different model reduction techniques have been proposed, e.g. sequential elimination of insignificant terms or selecting between only full VAR models with different maximum lag orders, among other. I have chosen two techniques: simplified all subset selection and elastic net regularization employed on a full VAR(5) model.

3.2.5.1 Full subset selection

In all subset selection, the starting point is a full VAR(p) model containing $2^{K \cdot K \cdot p}$ different subset models made up of $K \cdot 2^{K \cdot p}$ unique equations. The goal is to select the subset model with the lowest AIC value. Estimating all the different subset models for obtaining their AIC values amounts to running $K \cdot 2^{K \cdot p}$ OLS regressions in the first stage of the FGLS estimation and $2^{K \cdot K \cdot p}$ second-stage FGLS regressions. Some simplifications are necessary to reduce the computational burden to an acceptable level. FGLS estimation could be replaced by equation-by-equation OLS estimation to trade off estimation efficiency against computational gains. An efficient branch-and-bound algorithm (Lumley and Miller, 2009) could be used to select the AIC-minimizing subset equation for each of the individual equations of the full VAR model. Assuming that innovations across the model equations are independent would allow adding up the individual equations' AIC values to obtain the AIC of the whole VAR model; hence, the subset equations that individually minimize the AIC would constitute a VAR model that minimizes the AIC across all the subset VAR models. Even though this assumption will likely be violated in practice, the gains in computational time are decisive. Thus I proceed as follows. For each equation of the full VAR model I find the subset equation that has the smallest AIC value. I combine the subset equations together to construct the VAR model to be used for forecasting.

The problem of over-fitting may be especially severe when using the all subset selection. For one full equation of a VAR model with five variables and five lags, the number of subset equations with the same number of regressors can be as large as $\binom{K \cdot p}{int(K \cdot p/2)} \approx 5 \cdot 10^6$. AIC-based selection between these subset equations will be equivalent to selecting the likelihood-maximizing model and will effectively be a form of data dredging. As a means to alleviate (although not entirely cure) the problem, a number of models with low AIC values could be considered instead of one, and their forecasts averaged. For each model equation, I consider up to 40 lowest-AIC subset equations to form a subset VAR model. That is, I form one VAR model from the set of subset equations that each have the lowest AIC within all subsets; then I form another VAR model from the set of subset equations that each have the second-lowest AIC within all subsets; and so on up to the 40th-lowest AIC. I discard the models whose AIC values are two or more units above the AIC of the first model; the constant 2 is selected due to a rule of thumb given in Burnham and Anderson (2004).

3.2.5.2 Regularization

The second approach towards lag order selection and model reduction is regularization. It may be used broadly whenever the goal is to lessen the number of independent variables or limit their contributions to the dependent variable. The idea of regularization is to reduce the model variance by shrinking the estimated model coefficients towards zero. This comes at an expense of increased model bias. However, if we start from a relatively richly parameterized model, the reduction in the model variance due to shrinkage will likely outweigh the increase in the squared model bias. This will lead to a decrease in the mean squared forecast error. Naturally, the shrinkage should stop when the squared model bias starts increasing fast enough to outweigh the reduction in the model variance. The optimal amount of shrinkage can be found using cross validation. See Hastie et al. (2009, p. 61, 241-249) for details.

Regularization can be done in a few different ways. For example, imposing a penalty term on the sum of absolute values of model coefficients during estimation is known as L^1 regularization, or the LASSO (Tibshirani, 1996). Penalizing the sum of squared values of the model coefficients is L^2 regularization and is used in the ridge regression (Hoerl and Kennard, 1970). A combination of the LASSO and the ridge regression is called the elastic net regression (Zou and Hastie, 2005). The different regularization techniques suit different problems; it is not always clear whether to choose the LASSO, the ridge regression or the elastic net regression. The benefit of the elastic net is that it is flexible; cross validation can be used to select the optimal relative weights for the L^1 and L^2 penalties. I do that by searching over a grid of weights between $[0,1]$ and $[1,0]$ with a step of 0.1. Given the relative weights, the optimal amount of shrinkage is selected by searching over a grid of shrinkage parameter values and using the so-called one-standard-error rule. That is, I aim for the maximum shrinkage as long as the training error remains within one standard deviation of the minimal training error across the grid of all the shrinkage values (Hastie et al., 2009, p. 244). I apply regularization separately for each equation of the full VAR(5) model. I also choose the relative weights for the L^1 and the L^2 penalties and the shrinkage parameters independently across the equations. Out of the two model reduction techniques, I expect regularization to be more fruitful in forecasting than the simplified all subset selection since regularization is less susceptible to over-fitting.

3.2.6 Vector error correction model

A special case of the VAR model used for cointegrated variables is the vector error correction model (VEC model, or VECM). It can be represented either as a VAR model with nonlinear coefficient restrictions for the variables in levels or as a VAR model appended by a so-called error correction term for the variables in their first differences. The latter is the more common representation and is given below:

$$\Delta \boldsymbol{y}_t = \boldsymbol{c} + \alpha(\beta' \boldsymbol{y}_{t-1}) + \Gamma_1 \Delta \boldsymbol{y}_{t-1} + \cdots + \Gamma_p \Delta \boldsymbol{y}_{t-p} + \boldsymbol{\varepsilon}_t;$$

α and β are rectangular $K \times r$ matrices where r denotes the cointegration rank; $\beta' \boldsymbol{y}_{t-1}$ is a vector consisting of stationary combinations of the variables in levels, known as the error correction term. If the system of variables evolves according to a VECM, neither an unrestricted VAR in levels nor an unrestricted VAR in the first differences is a well specified model. The unrestricted VAR in levels fails to incorporate the non-linear restrictions due to the presence of cointegration; the VAR in the first differences suffers from an omitted variable bias due to the absence of the error correction term. Since the VECM has more parameters than either the VAR in levels or the VAR in the first differences, the model estimation variance tends to be relatively larger, which yields increased forecast errors; see Hastie (2009, p. 223-228) for a general treatment. In practice, evidence is mixed as to whether VECMs yield better forecasts than unrestricted VAR models for short forecast horizons; see e.g. Chigira and Yamamoto (2009), Christoffersen and Diebold (1998), Engle and Yoo (1987), Hoffman and Rasche (1996), and Löf and Franses (2001).

To make the VECM model operational, first the cointegration rank of the system of variables needs to be determined. I examine the potential cointegration between the spot price, the share price index and the futures predictive variable in the initial subsample covering 2007w1-2012w52 using the standard Johansen procedure (Johansen, 1991). One cointegrating vector, or equivalently two stochastic trends, is found. That is, the spot price, the share price index and the futures predictive variable share two long-run trends. Although one stochastic trend among the spot price, the share price index and the front month futures series may seem more natural, recall that the futures predictive variable may diverge from the front month futures by construction, bringing in the second stochastic trend. The presence of only one cointegrating vector allows obtaining the error correction term using the first step of the Engle-Granger two-step procedure (Engle and Granger, 1987). The spot price is regressed on the other two variables, and the regression residual becomes the error correction term. Cointegration analysis by the Johansen procedure applied on all 92 6-year-rolling- window subsamples yields zero cointegrating vectors in 52 cases and one cointegrating vector in 40 cases. Zero cointegrating vectors implies no cointegration. If the VECM is used regardless, the error correction terms are redundant. Including a redundant variable in a linear model should increase the model variance but should not affect the consistency of the parameter estimates. Both the simplified all subset selection and the elastic net regularization should minimize the impact of a redundant variable in the system. Thus I stick to forcing one cointegrating vector and using one error correction term obtained as described above. Imposing the number of cointegrating vectors to be one simplifies the automation of forecasting by the VECM.

The lag order selection problem for the VECM is analogous to that for the VAR model. I treat it using the simplified all subset selection and the elastic net regularization just as in the case of the VAR model above.

3.2.7 Artificial neural network

Artificial neural networks (ANNs) are a class of statistical learning methods used to approximate relationships of unknown functional form among a large number of variables. The independent variables are called inputs, and the dependent variable is called output in the ANN terminology. The output from a network is a non-linear function of a weighted sum of the inputs. The ANNs are flexible enough to approximate any continuous function arbitrarily well (Hastie et al., 2009, p. 390). However, this flexibility comes at a cost; the ANNs are prone to over-fitting and they are also computationally intensive. I will use a relatively simple one-hidden-layer feed-forward network as depicted in Figure 9. It consists of K+1 input nodes (K regressors and a constant term), one output node and one intermediate layer of m hidden nodes.

Each hidden node i works as follows: first it produces a weighted sum of all the inputs; second, it transforms the weighted sum using the sigmoid function: $\sigma(x) = \frac{1}{1+\exp{(x)}}$. Hence, the output from the node i is

$$z_i = \frac{1}{1 + \exp{(w_{i0} + w_{i1}x_1 + \cdots + w_{ik}x_k)}}$$

where x's are the inputs and w_i's are weights specific to the node i. Meanwhile, the output node does half of the job of a hidden node; it takes in outputs from all the hidden nodes and produces a linear combination of their values plus a constant term:

$$y = v_0 + v_1 z_1 + \cdots + v_m z_m;$$

here vs are weights specific to the output node. Combining the effects of the hidden nodes and the output node, we obtain the output y in terms of the inputs x_1, …, x_k:

$$y = v_0 + v_1 z_1 + \cdots + v_m z_m = v_0 + \frac{v_1}{1+\exp(w_{10}+w_{11}x_1+\cdots+w_{1k}x_k)} + \cdots + \frac{v_m}{1+\exp(w_{m0}+w_{m1}x_1+\cdots+w_{mk}x_k)}.$$

Fitting an ANN amounts to determining the weights ws and vs. There is no algebraic solution to the ANN problem, thus optimization is used. Weights are set at their initial values at the beginning and then iteratively updated in the direction that gives the largest improvement in the model fit. Due to the flexibility of the ANN, iterating until convergence will likely lead to over-fitting. Two solutions have been proposed to tackle this issue. A historically popular approach known as early stopping restricts the number of iterations (Hastie et al., 2009, p. 398). This saves computational time at the same time as it prevents over-fitting. Alternatively, regularization can be used by penalizing the total sum of absolute or squared weights, similarly to the LASSO or the ridge regression. I choose to combine both approaches. I limit the number of iterations to 100, the default value in the "nnet" package in "R" (Venables and Ripley, 2002); I also penalize the sum of the squared weights by setting the penalty parameter to 0.1 as recommended in Hyndman and Athanasopoulos (2014).

Figure 9 Schematic representation of a one-hidden-layer feed-forward artificial neural network

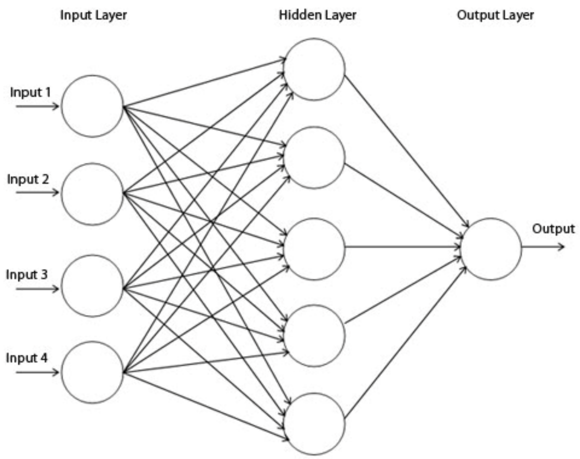

[11]

Since optimization is used for fitting the ANN, the choice of the starting values for the weights is relevant. Hastie et al. (2009, p. 397-398) advise to set the starting values close to zero where the sigmoid function is approximately linear. The iterative optimization moves the weights possibly away from zero to the area of greater nonlinearity. Thus the ANN starts as a roughly linear model and allows for introducing nonlinearity during the optimization. Following the recommendations for starting weights in the "nnet" package in "R" (Venables and Ripley, 2002), I use random uniform weights over the interval $[-d, d]$ where d is the inverse of the maximum absolute value across all the inputs, i.e. $d = \frac{1}{\max_{1 \le i \le k}(\max_t (|x_{i,t}|))}$.

Different starting values will generally produce different final weights ws and vs, especially when convergence may not be achieved due to the early stopping. I follow the common practice (Hyndman and Athanasopoulos, 2014) of estimating a number of models differing only by the starting values, and average the final result. I use 20 different sets of starting values, the number being limited by the computational cost.

An important element of building ANNs is the selection of the number m of nodes in the hidden layer. The existing rules of thumb are diverse and sometimes contradictory, thus I do the selection by cross validation. I try out ANNs with $m = 1, \dots, 20$ hidden nodes and select the m that produces the lowest RMSE in the validation sample. Computational costs are again a concern; the greater the number of nodes, the longer it takes to fit the ANN. When applied to the data I have, no particular

[11] The picture is borrowed from http://www.codeproject.com/Articles/175777/Financial-predictor-via-neural-network

hidden layer size dominates in cross validation, and the selected size varies considerably across the rolling windows.

Direct forecasts seem more natural than iterated ones when forecasting with the ANNs. Therefore, I build a separate ANN for each forecast horizon.

3.2.8 K-nearest neighbours method

K-nearest neighbours (kNN) method is a non-parametric pattern recognition technique used for classification and regression. It is based on the idea that similar situations and similar sequences of events can be encountered repeatedly over time. If we have observed a sequence A, B, C, D multiple times in the past and are now seeing A, B, C, we might expect D to follow. That is, having observed A, B, C we can *forecast* the next observation to be D. Of course, we need not find exact, mechanic repetitions of identical sequences in economic time series, but the idea can be adapted to similar rather than identical situations.

Let us define a "situation" as a time sequence of observations of a group of variables; in other words, a "situation" is a short multivariate time series. In our application, the variables are the spot price, the export volume, the share price index, the EUR/NOK exchange rate and the futures predictive variable, and their lags. The lag length reflects how long a history is relevant when comparing the "situations". I allow the lag length to vary between two and nine and choose the optimal lag length by cross validation.

The similarity between any two "situations" can be measured by a distance function. The distance function should be defined so as to produce small values for similar "situations" and large values for dissimilar ones. I use Euclidean distance for simplicity although other distance functions could be considered as well. The Euclidean distance varies with both the location and the scale of the data, which may or may not be useful when comparing the "situations". If the "situations" were normalized with respect to location, only the cumulative changes in the variables over the lag length would matter, and the starting level would be disregarded. For example, this would make the distance measure indifferent as to whether the spot price is rising from a relatively low level or whether it is rising from an already high level and breaking a record. I assume these two situations are quite different in the perspective of the salmon spot market, thus I do not adjust for location. Meanwhile, adjustment for scale would correspond to viewing fluctuations relative to their typical level. For example, an increase in export volume from 9 000 to 10 000 tonnes on a slow week may have a different impact on the spot price than an increase from 19 000 to 20 000 tonnes on a busy week, which makes scale-adjusted distance relevant. Therefore, I choose to adjust the data for scale when calculating the distance. Admittedly, both choices are somewhat arbitrary and could be debated. Meanwhile, considering only the (seasonally-adjusted) log-levels of the data makes the nearest neighbours almost always be the adjacent observations in time, which is not desirable. Including the log-returns extra to the log-levels in the definition of a “situation” yields a noticeable improvement.

Once a “situation” and the distance function are defined, the number of relevant neighbours, k, needs to be determined. It could be fixed, say, five or ten; or it could depend on the distance so that every neighbour within a given distance from the current “situation” would be considered. I choose the first option due to my sample size being quite small; in large samples, the second option seems more promising. I further choose the number k of relevant neighbours from between five and ten using cross validation.

Given the neighbours and the “situation” of interest, a forecast needs to be constructed. Two different approaches may be used. First, one could ignore the data from the current “situation” and obtain the forecast using only the information contained in the neighbour “situations”. For example, a (weighted) average of the one-step-ahead spot prices of the neighbours could serve as the one-step-ahead forecast of the current spot price. Note that such a prediction would never end up outside the range of the neighbours’ values, which seems undesirable; this can be partly circumvented by using scaled data as I do. Second, one could use the data from the current “situation”, too. That would amount to running a locally-weighted regression where only the neighbours (instead of all the data points) would be used to estimate an assumed functional relationship, e.g. a local VAR model. The locally-estimated model would be employed for predicting one or multiple steps ahead from the current “situation”. The latter approach is attractive but requires a rather large sample to work reliably. Therefore, my choice is the former approach.

Having discussed the individual models used in this study, I will now consider combinations of their forecasts before turning to the empirical results.

3.3 Forecast combinations

When the form of the data generating process is unknown, different models may be applied to approximate it. Even the best individual approximation need not be better than all the remaining approximations at all times. Thus it is natural to use different model at different times. For example, if a process has two regimes and the first one is well approximated by an ARIMA model while the second is better approximated by a neural network, we would wish to use different models for the different regimes. However, in practical applications we may not be able to clearly identify the regimes and distinguish when a particular model is to be preferred over its competitors. Under such circumstances, a linear combination of forecasts from a few different models may be employed. Forecast combinations often outperform individual models’ forecasts and even the forecasts of the best of the individual models; this is known as the forecast combination puzzle (Timmermann, 2006).

There are a number of ways to combine forecasts from different models. A simple average forecast weighs all the candidate forecasts equally and is very easy to implement. In a trimmed mean method, the smallest and the largest forecasts are discarded and the mean is taken over the remainder. In a winsorized mean approach, the smallest and the largest forecasts are replaced by the smallest and the largest values of the trimmed set of forecasts. More sophisticated schemes give larger weights to individual forecasts from the models with relatively low AIC or BIC values. The relative performance of the different weighting schemes has been mixed in practice, and the simple combinations often perform as well as, or better than, the more complicated ones (Timmermann, 2006); see Diebold (2014, p. 471-481) and Smith and Wallis (2009) for explanations. Due to this reason and for the sake of simplicity only the simple unweighted average and the trimmed mean approach will be used in this study.

3.4 Performance assessment

The forecasting performance of a given forecast can be assessed from a number of perspectives. First, forecast accuracy can be measured in terms of how close (in absolute or percentage terms) the forecast is to the realized value on average. Second, forecast accuracy may be compared to that of a benchmark forecast produced by some naïve model. Third, forecast optimality can be investigated by examining whether the forecast errors are unpredictable. Fourth, competing forecasts can be compared by inspecting the losses produced by the different forecast errors, or by trying to find one forecast beyond which the competing forecasts do not help reduce the forecast error significantly.

3.4.1 Measures of forecast accuracy

The forecast accuracy measures used in this paper are the following: root mean squared error, mean absolute error, mean absolute percentage error, mean absolute scaled error, Theil's U statistic and hit rate. Reporting a number of different accuracy measures is intended for facilitating evaluation of forecasting performance using different loss functions; the latter may depend on the intended use of the forecasts. The first accuracy measure is root mean squared error (RMSE), defined as

$$RMSE = \sqrt{\frac{1}{T_1 - h - T_0 + 1}\sum_{t=T_0}^{T_1-h}\left(f_{t,h} - r_{t+h}\right)^2}$$

where $f_{t,h}$ is the h-step ahead forecast as of time t; r_{t+h} is the realized value at time $t+h$; and T_0 and T_1 are the first and the last periods of the "hold-out sample", respectively. It is a natural error measure under squared loss. RMSE is the same regardless of whether the forecast and the realized value is the price change or the price level. The same is true for mean absolute error (MAE) which is defined as

$$MAE = \frac{1}{T_1 - h - T_0 + 1} \sum_{t=T_0}^{T_1 - h} \left| f_{t,h} - r_{t+h} \right|.$$

MAE matches a loss function that is linear in the size of the error and indifferent to its sign. Mean absolute percentage error (MAPE) is

$$MAPE = \frac{1}{T_1 - h - T_0 + 1} \sum_{t=T_0}^{T_1 - h} \frac{\left| f_{t,h} - r_{t+h} \right|}{r_{t+h}}.$$

It is natural to define MAPE over the forecast and the realized value of the price level rather than the price change. MAPE treats errors of the same absolute size differently depending on the reference level. E.g. an error of $\pm$NOK 1/kg is punished harder when the price level is NOK 30/kg than when it is NOK 40/kg; NOK 1/kg constitutes a larger percentage of the former level than the latter one.

Mean absolute scaled error (MASE) due to Hyndman and Koehler (2006) is defined as

$$MASE = \frac{1}{T_1 - h - T_0 + 1} \sum_{t=T_0}^{T_1 - h} \frac{\left| f_{t,h} - r_{t+h} \right|}{d}$$

where $d = \frac{1}{T_0 - h} \sum_{t=1}^{T_0 - h} \left| f_{t,h} - r_{t+h} \right|$ is the "in-sample" mean forecast error produced by a naïve no-change forecast. It compares the MAE of the actual forecast in the "hold-out sample" against the "in-sample" MAE of a naïve no-change forecast. A reading below unity indicates that the actual forecast does better out of sample than a naïve no-change forecast does in sample; a value above unity suggests the opposite.

Theil's U statistic, defined as

$$U = \frac{\sqrt{\frac{1}{T_1 - h - T_0 + 1} \sum_{t=T_0}^{T_1 - h} \left(\frac{r_{t+h} - f_{t,h}}{r_{t+h}} \right)^2}}{\sqrt{\frac{1}{T_1 - h - T_0 + 1} \sum_{t=T_0}^{T_1 - h} \left(\frac{r_{t+h} - r_t}{r_{t+h}} \right)^2}}$$

is another measure that compares the forecasts of interest to a naïve no-change forecast, both measured in the "hold-out sample". Note that there is considerable confusion as Theil has defined the U statistic in two different ways (Bliemel, 1973) and then other sources have defined it in several new ways that are different from any of the two original definitions; see e.g. Diebold (2014, p. 440) and Makridakis and Hibon (1995). I choose the widespread definition of Makridakis and Hibon (1995).

Hit rate, defined as

$$Hit\ rate = \frac{1}{T_1 - h - T_0 + 1} \sum_{t=T_0}^{T_1 - h} I\left\{ sign\left(f_{t,h} - r_t \right) = sign\left(r_{t+h} - r_t \right) \right\}$$

is the proportion of the times when the forecasted price change has the same sign as the realized one. The sign of the price change is relevant whenever salmon can be stored in the sea at no additional cost for another week. Then it pays to slaughter the fish in the week with the higher spot price. Although the zero cost assumption is quite strict, it can be realistic at certain times of year when the salmon is neither fed nor vaccinated, nor otherwise maintained.

3.4.2 Tests of forecast optimality and encompassing

Diebold (2014, p. 403-441) lists a number of forecast optimality criteria. The unifying principle is that the forecast errors should be unpredictable using the information available at the time of making the forecast. Otherwise, the forecasts could be improved by exploiting the predictable component of the forecast errors; thus such forecasts would not be optimal. There are four conditions that the forecast errors of an optimal forecast satisfies. First, an optimal forecast is unbiased; hence, the forecast errors have a mean of zero. Second, one-step-ahead forecast errors are white noise. Third, h-step-ahead forecasts are at most an MA($h-1$) process. Fourth, variances of h-step-ahead forecast errors are non-decreasing in h.

The first condition can be checked by estimating the mean of the forecast errors and testing whether it is significantly different from zero. The second condition can never be fully evaluated due to the sheer variety of the potential forms of dependence. However, at least one simple check could be performed. Since white noise implies no autocorrelations, the Ljung-Box test (Ljung and Box, 1978) could be used to assess whether the error series are autocorrelated. The third condition can be tested by fitting an MA($h-1$) model to the forecast errors and examining the model's adequacy by running the Ljung-Box test on the model residuals. Finally, the fourth condition can be checked by ranking the empirical error variances.

The four optimality conditions listed above consider only the information contained in the forecast errors. Expanding the information set allows conducting another test of forecast optimality. Under forecast optimality, a regression of the forecast errors on the forecasts themselves should yield a zero intercept and a zero slope coefficient. A non-zero intercept and/or slope signals that the information has not been optimally utilized when constructing the forecasts. Equivalently, one could regress the realized values on the forecasts and test for a zero intercept and a unit slope. This is called the Mincer-Zarnowitz regression and the Mincer-Zarnowitz test (Mincer and Zarnowitz, 1969).

I will run the test on price changes in place of price levels for the following two reasons. First, forecasts of price levels will tend to track the realized values closely by construction, since the forecast origin is the actual (realized) price as of h periods before the target period (except for the case of the kNN forecasts). Therefore, even relatively poor forecasts will tend to have a unit slope and perhaps also a zero intercept. Second, it will help avoid potential modelling issues due to the realized values and the forecasts likely being integrated and cointegrated in levels. Also, h-step-ahead forecasts may be expected to have serial correlation as the forecasts overlap at horizons greater than one, which may result in serially correlated model errors. I will thus use heteroskedasticity-and-autocorrelation-consistent (HAC) covariance matrix estimator (Newey and West, 1987) when obtaining the Mincer-Zarnowitz test statistic.

The relative performance of any pair of forecasts can be gauged using the Diebold-Mariano test (Diebold and Mariano, 1995; see also Diebold, 2015). The test assesses whether the loss due to the forecast errors is equally great for both forecasts. The forecast errors from the two forecasts are transformed into corresponding losses using a loss function, e.g. by taking the absolute values of the errors. Then a loss differential is constructed as the difference between the two loss series. If the two forecasts are equally good in terms of the chosen loss function, the population mean of the loss differential series should be equal to zero. The Diebold-Mariano test amounts to checking the latter hypothesis.

When multiple forecasts are available and one of them is taken as a benchmark, the relevance of all the other forecasts could be questioned. A benchmark forecast is said to encompass the other forecasts if given the benchmark forecast, the competing forecasts do not reduce the forecast error. This may be investigated by testing the null hypothesis of $\beta_i = 1, \beta_m = 0$ for $m \neq i$ in the regression

$$r_{t+h} = \beta_1 f_{t,h}^1 + \cdots + \beta_m f_{t,h}^M + \varepsilon_{t+h} \qquad (1)$$

where the i^{th} forecast $f_{t,h}^i$ is taken as the benchmark. If the null hypothesis cannot be rejected, the benchmark forecast encompasses the other forecasts. If none of the forecasts encompasses the others, forecast combinations are likely to be superior to any individual forecast (Diebold, 2014, p. 463-464). I will employ the test on price changes rather than levels and will use HAC covariance matrix estimator by the same argumentation as for the Mincer-Zarnowitz test. This approach is also discussed in Newbold and Harvey (2008).

4 Results

Table 2 displays six forecast accuracy measures for all the different forecasting methods and forecast horizons. The corresponding Figure 10 illustrates two of the accuracy measures, MAE and hit rate. A look at the naïve no-change or the naïve seasonal forecasts shows that a naïve forecast may be hard to beat. The naïve seasonal forecast performs relatively well with respect to all of the accuracy measures and delivers the best results in nine cases out of 30 (six error measures times five forecast horizons). Among the individual forecasts, a clear leader is the VECM estimated using the elastic net regularization (VECM-EN). It delivers good results regardless of the accuracy measure and has 14 best forecasts out of 30. A definite loser is the ANN forecast which does considerably worse than any other competitor. (Changing the weight decay parameter or decreasing the number of nodes to prevent over-fitting did not improve the ANN's forecasting performance.) The two forecast combination methods do not generally beat the individual forecasts except for one best result out

Table 2 Forecast accuracy

Accuracy measure	Horizon	No change	Seasonal	ARIMA	ARFIMA	ARARMA	VAR 1 best	VAR n best	VAR-EN	VECM 1 best	VECM n best	VECM-EN	kNN	ANN	Mean forecast	Trimmed mean
RMSE	1	2,288	2,257	2,431	2,441	2,431	2,318	2,305	2,364	2,313	2,307	2,248	***2,225**	**3,308**	2,285	2,258
	2	3,603	3,546	3,772	3,867	3,772	3,716	3,673	3,693	3,654	3,622	***3,523**	3,658	**4,427**	3,641	3,644
	3	4,425	4,304	4,381	4,483	4,381	4,399	4,363	4,471	4,330	4,291	***4,213**	4,669	**5,098**	4,339	4,338
	4	5,079	4,838	4,945	4,929	4,945	4,927	4,904	5,100	4,816	***4,801**	4,805	5,018	**5,727**	4,816	4,816
	5	5,475	***5,098**	5,249	5,291	5,249	5,326	5,326	5,444	5,237	5,199	5,149	5,669	**6,171**	5,210	5,202
MAE	1	1,771	1,753	1,937	1,919	1,937	1,874	1,875	1,917	1,897	1,878	1,816	***1,697**	**2,702**	1,845	1,828
	2	2,904	2,882	3,038	3,153	3,038	3,029	3,004	3,007	2,968	2,937	***2,757**	2,850	**3,617**	2,961	2,953
	3	3,448	3,488	3,572	3,623	3,572	3,645	3,604	3,586	3,585	3,560	***3,444**	3,744	**4,300**	3,551	3,532
	4	3,886	***3,845**	3,989	4,011	3,989	3,970	3,960	4,172	3,879	3,861	3,908	3,970	**4,733**	3,900	3,917
	5	4,282	***4,056**	4,348	4,339	4,348	4,399	4,402	4,509	4,318	4,287	4,261	4,669	**4,976**	4,307	4,319
MAPE	1	4,536	4,518	4,968	4,923	4,968	4,859	4,862	4,954	4,897	4,852	4,687	***4,406**	**6,992**	4,769	4,734
	2	7,412	7,383	7,755	7,985	7,755	7,791	7,730	7,728	7,607	7,540	***6,990**	7,360	**9,183**	7,593	7,577
	3	8,871	8,933	9,076	9,131	9,076	9,278	9,180	9,201	9,089	9,041	***8,698**	9,623	**10,756**	9,014	8,975
	4	9,968	9,793	10,086	9,992	10,086	10,034	10,005	10,626	9,800	***9,764**	9,813	10,083	**11,700**	9,831	9,884
	5	10,991	***10,260**	10,951	10,741	10,951	11,048	11,044	11,415	10,839	10,762	10,636	11,845	**12,206**	10,792	10,842
MASE	1	0,985	0,976	1,075	1,067	1,075	1,040	1,040	1,062	1,055	1,044	1,010	***0,945**	**1,498**	1,023	1,014
	2	1,181	1,173	1,231	1,279	1,231	1,228	1,218	1,218	1,206	1,193	***1,120**	1,158	**1,463**	1,200	1,197
	3	1,277	1,292	1,319	1,340	1,319	1,346	1,331	1,324	1,328	1,319	***1,275**	1,388	**1,587**	1,312	1,305
	4	1,390	***1,373**	1,422	1,434	1,422	1,415	1,412	1,485	1,386	1,380	1,397	1,417	**1,690**	1,390	1,396
	5	1,413	***1,333**	1,429	1,432	1,429	1,446	1,447	1,477	1,425	1,414	1,406	1,533	**1,644**	1,416	1,419
Theil's U	1	1,000	***0,995**	1,061	1,068	1,061	1,034	1,030	1,061	1,025	1,024	***0,995**	***0,995**	**1,464**	1,019	1,007
	2	1,000	0,980	1,038	1,054	1,038	1,039	1,029	1,034	1,017	1,011	***0,964**	1,041	**1,221**	1,015	1,016
	3	1,000	0,962	0,971	0,975	0,971	0,974	0,969	1,001	0,957	0,950	***0,920**	1,047	**1,093**	0,958	0,958
	4	1,000	0,942	0,948	0,915	0,948	0,947	0,943	0,993	0,923	0,922	***0,912**	0,970	**1,051**	0,917	0,919
	5	1,000	0,915	0,929	0,902	0,929	0,939	0,939	0,975	0,926	0,920	***0,901**	1,023	**1,042**	0,915	0,916
Hit rate	1	---	***0,648**	0,593	0,527	0,593	0,571	0,560	0,604	0,571	0,560	0,505	0,637	**0,462**	0,560	0,571
	2	---	0,600	***0,644**	**0,556**	***0,644**	0,611	0,622	0,622	0,633	0,633	***0,644**	0,600	0,578	0,622	0,622
	3	---	0,539	0,584	***0,607**	0,584	0,573	**0,528**	0,562	0,584	***0,607**	0,573	0,562	0,596	***0,607**	0,584
	4	---	0,534	**0,511**	0,568	**0,511**	0,523	0,580	0,545	***0,614**	0,602	0,557	0,568	0,557	0,557	0,534
	5	---	***0,644**	0,552	0,586	0,552	**0,529**	0,540	0,552	0,575	0,575	0,575	0,563	0,552	0,540	0,563

Note: forecasts in bold with an asterisk are the best across the methods for the particular forecast horizon and the particular accuracy measure (i.e. the best across the table row); forecasts in bold without an asterisk are the worst. VAR 1 best is the best VAR model from all subset selection; VAR n best is the average forecast of the n best models from all subset selection, where the choice of n is described in the text; VAR-EN is the VAR model estimated using elastic net regression; VECM 1 best, VECM n best and VECM-EN are the corresponding VECM models. ARARMA forecasts are identical to ARIMA forecasts for this particular dataset.

Figure 10 Forecast accuracy measures for forecast horizons 1 through 5

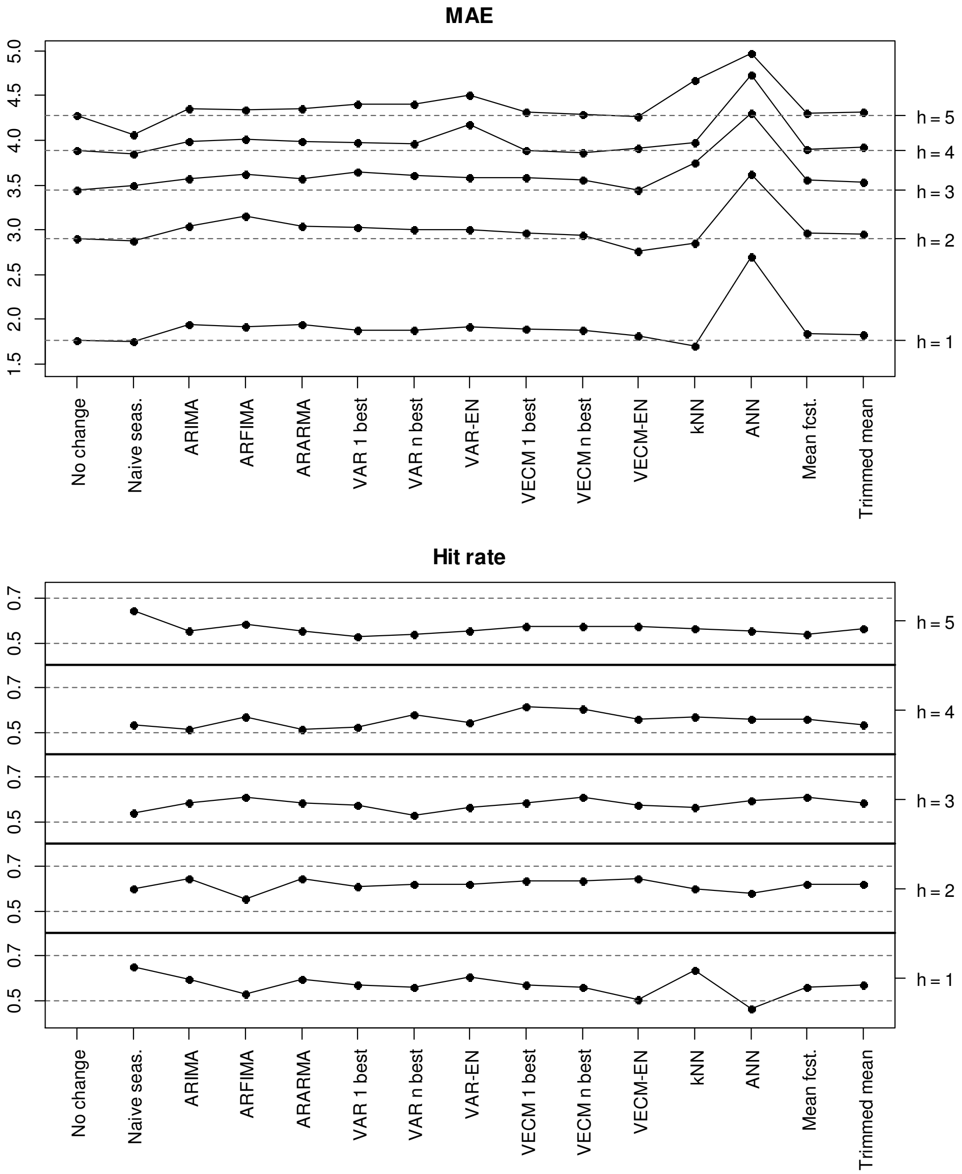

Note: on the right axis, $h = 1$ through $h = 5$ index the forecast horizons. Lower MAE is better. Dashed horizontal lines in the MAE plot mark the MAE of the naïve no-change forecast. Higher hit rate is better. Hit rate below 0.5 is worse than a random guess. VAR 1 best is the best VAR model from all subset selection; VAR n best is the average forecast of the n best models from all subset selection, where the choice of n is described in the text; VAR-EN is the VAR model estimated using elastic net regression; VECM 1 best, VECM n best and VECM-EN are the corresponding VECM models.

of 30. Only two methods – the naïve seasonal and the VECM-EN – beat the naïve no-change forecast for all the five forecast horizons as measured by Theil's U statistic.

The forecast accuracy ranks averaged over the different horizons are given in Table 3. Again, the VECM-EN forecast fares the best with the top rankings in terms of RMSE, Theil's U and overall; the naïve seasonal forecast comes in second in the overall ranking having topped MAE, MAPE and MASE rankings. The ANN forecast comes in last for all accuracy measures but hit rate, where the naïve no-change forecast is formally the worst since it always predicts zero change. Having discovered that the VECM-EN and the naïve seasonal are the two best forecasting methods, I will focus on them in more detail. The naïve seasonal forecasts together with the realized values are depicted in Figure 11.

I will examine the optimality of the two forecasts. Unbiasedness cannot be rejected for either of the forecasts; the test's p-values are all above 0.50. I carry out the Ljung-Box test for lags one through 20 to assess the statistical significance of autocorrelations in the one-step-ahead forecasts. Numerous autocorrelations at nearby lags are statistically significant for the naïve seasonal forecast; a significant autocorrelation at lag one is found for the VECM-EN forecast. Therefore, the hypothesis that one-step-ahead forecasts are white noise is rejected.

To test whether the h-step-ahead forecast errors are at most MA($h-1$) processes for $h = 2, \dots, 5$, I first fit an MA($h-1$) model to the forecast errors and then examine the model's residuals. Statistically significant autocorrelations at nearby lags are found for all the forecast horizons in case of the naïve seasonal forecasts, as evidenced by the Ljung-Box test. Meanwhile, none of the residual autocorrelations are statistically significant for the VECM-EN forecast errors, for any of the forecast horizons. Thus, the naïve seasonal forecast fails this optimality test while the VECM-EN passes it.

To see whether the forecast error variances are increasing in the forecast horizon h, I examine the empirical variances for both the naïve seasonal and the VECM-EN forecasts. Indeed, the error variances are increasing in both cases.

In sum, the naïve seasonal forecast meets two out of four optimality conditions while the VECM-EN fulfils three out of four.

The Mincer-Zarnowitz test results indicate non-optimality for both the naive naïve seasonal and the VECM-EN forecasts. The null hypothesis of optimality is strongly rejected for all the forecast horizons.

Considering the Diebold-Mariano test, a loss function associated with the forecast error has to be defined first. I choose absolute loss since salmon farmers' revenues or processors' costs depend linearly on the salmon price. I will use the Diebold-Mariano test to compare the naïve no-change forecast, the naïve seasonal forecast and the VECM-EN forecast against each other. Modified version of the original test statistic due to Harvey et al. (1997) will be employed to account for

Table 3 Forecast accuracy rankings (averages over forecast horizons)

Accuracy measure	Horizon	No change	Seasonal	ARIMA	ARFIMA	ARARMA	VAR 1 best	VAR n best	VAR-EN	VECM 1 best	VECM n best	VECM-EN	kNN	ANN	Mean forecast	Trimmed mean
RMSE	All	9,2	3,0	10,3	11,8	10,3	9,9	8,1	11,8	6,0	3,6	***1,6**	9,8	**15,0**	5,0	4,6
MAE	All	3,2	***2,0**	10,9	11,8	10,9	10,3	9,6	11,6	7,2	5,2	2,8	8,1	**15,0**	5,6	5,8
MAPE	All	5,2	***2,2**	10,7	9,2	10,7	11,2	10,0	12,0	7,2	4,6	2,4	8,4	**15,0**	5,6	5,6
MASE	All	3,3	***2,0**	10,6	12,2	10,6	10,3	9,6	11,3	7,6	5,4	3,0	8,2	**15,0**	5,5	5,4
Theil's U	All	9,2	4,1	10,0	8,6	10,0	10,3	8,5	11,8	6,2	4,8	***1,2**	11,0	**15,0**	4,4	4,9
Hit rate	All	**15,0**	7,4	7,2	6,9	7,2	10,5	9,4	8,1	4,6	***4,5**	7,1	7,2	9,5	7,8	7,6
Average	All	7,5	3,5	10,0	10,1	10,0	10,4	9,2	11,1	6,5	4,7	***3,0**	8,8	**14,1**	5,7	5,7
Overall	All	7,0	2,0	10,5	12,0	10,5	13,0	9,0	14,0	6,0	3,0	***1,0**	8,0	**15,0**	4,5	4,5

Note: lower rank is better. The naïve no-change forecast is ranked last with respect to hit rate because it always predicts the price change to be zero, which is never true in this dataset. VAR 1 best is the best VAR model from all subset selection; VAR n best is the average forecast of the n best models from all subset selection, where the choice of n is described in the text; VAR-EN is the VAR model estimated using elastic net regression; VECM 1 best, VECM n best and VECM-EN are the corresponding VECM models. ARARMA forecasts are identical to ARIMA forecasts for this particular dataset.

Figure 11 Naïve seasonal forecasts and realized values for forecast horizons 1 through 5

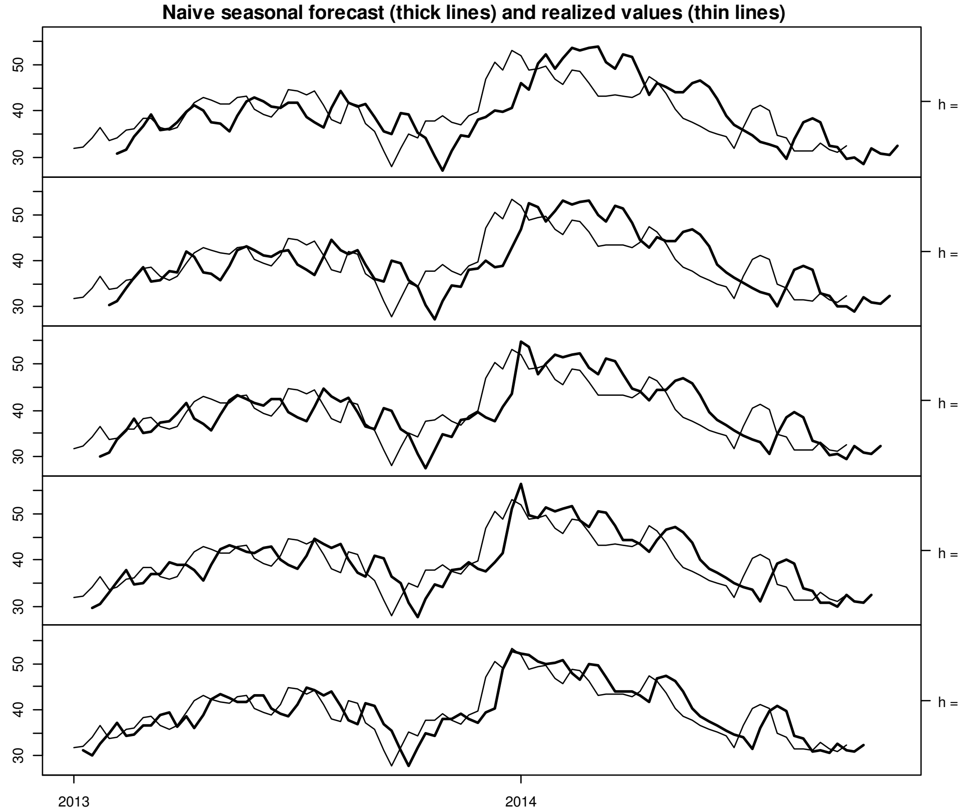

potential autocorrelations in the loss differentials. The loss differentials are found to be stationary in all cases. The test results are quite the same for all pairs of forecasts, for all the forecast horizons. Neither the naïve seasonal nor the VECM-EN forecasts improve upon the naïve no-change forecast when the loss function is the absolute value of the forecast error; p-values associated with the test statistic are always above 0.40. This is as expected because the naïve seasonal as well as the VECM-EN forecasts produce mean absolute errors that are close to the ones of the naïve no-change forecast; see Table 2. Apparently, the improvement of the naïve seasonal or the VECM-EN forecasts over the naïve no-change forecast is too small to be statistically significant. The economic significance is also quite small. The largest loss differential is found for the longest forecast horizon and averages at only NOK 0.23/kg. This is minor in light of the mean absolute error of the VECM-EN forecast being NOK 4.26/kg.

Next I examine forecast encompassing among the naïve no-change, the naïve seasonal and the VECM-EN forecasts. The test results suggest that neither of the forecasts encompasses the others, for any forecast horizon. Thus the forecasts are rather distinct despite having similar accuracy. As noted above, if none of the forecasts encompasses all the other ones, using a combination of forecasts could deliver improved accuracy. Although I have examined only the two best forecasts here, it seems unlikely that one of the other, generally inferior forecasts would encompass the rest. Still, forecast combinations do not perform exceptionally well as can be seen from their rankings in Table 3. The mean and the trimmed mean forecasts end up fourth and fifth in the overall ranking. That may be due to some of the inferior forecasts such as the ANN being highly inaccurate which spoils the performance of the forecast combinations.

Let us view the forecasting performance in light of the historical studies. None of the previous works on salmon price forecasting aimed at predicting the average (over the weight classes) spot price of salmon in the short run, so it is difficult to compare their results to mine. The closest is the article by Guttormsen (1999) who predicted weekly spot prices of disaggregated weight classes four, eight and twelve weeks ahead; only the four-week-ahead predictions are relevant here. Guttormsen (1999) explored only the spot price data (and no other variables) from 1992-1996 and presumably used an expanding window to estimate his models and predict the prices in 1997. I focus on the 4-5 kg weight class which is the closest to the weighted average spot price used in this study. A very low forecast error of 1.10% in terms of MAPE was achieved by the naïve no-change forecast, while a VAR(1) model delivered the second best forecast with MAPE of 1.21%. Compared with the MAPE of 9.97% for a naïve no-change forecast in my study, the results in Guttormsen (1999) look exceptionally good. However, the standard deviation of the salmon spot price in 1997 was as low as 0.86 compared to 5.60 in 2013w1-2014w39, a nearly sevenfold difference. Since the MAPE of the naïve no-change forecast in Guttormsen (1999) is about eight times lower than the MAPE of the corresponding forecast in this study, the relative forecast accuracy seems similar. Overall, Guttormsen (1999) concludes that the VAR(1) model delivered probably the best forecasts but still it could not outperform the naïve forecasts. The latter remark is echoed by the findings in this study; not even the best method consistently beats the naïve seasonal forecast, nor does it beat the naïve no-change forecast by a large margin.

5 Discussion

Accurate price forecasts are highly desirable in the volatile environment of the salmon spot market. They are instrumental in business planning for salmon farmers, processors, exporters and retailers, as well as futures market participants. While medium- and long-term price forecasts are supplied by commercial actors or could be inferred from the price quotes on the salmon futures contracts, short-term forecasts seem to be largely unavailable. This study is targeted to fill the existing gap.

I started out aiming to forecast the salmon spot price one to five weeks ahead. A bouquet of time series and machine learning methods were employed on historical data for 2007-2014 to generate pseudo out-of-sample forecasts. I used a six-year rolling window to select model specifications, estimate the models, and generate predictions. The forecasts were evaluated in terms of absolute accuracy, relative performance and optimality, and the different forecasting methods were ranked accordingly.

The following findings emerge from the study. First, all the forecasting methods get the direction of the price change right more than 50% of the time, for all forecast horizons[12], and three forecasting methods demonstrate at least 60% accuracy for the nearby two weeks. Knowing the direction of the price change is undoubtedly valuable as regards planning decisions of the market participants. The predictability in the salmon price may be explained by seasonality; indeed, a naïve-seasonal forecast is hard to beat when forecasting the price changes. Second, two forecasts, the naïve-seasonal and the VECM estimated using elastic net regularization (VECM-EN), beat the naïve no-change forecast consistently over the five forecast horizons as measured by Theil's U statistic. However, the gains are small from the economic perspective. Third, none of the methods delivers a considerable reduction in the mean absolute forecast error (MAE) as compared with a naïve no-change forecast. MAE is directly proportional to gains or losses in the salmon farmer's revenue or the processor's costs due to the variation in the spot price of salmon. Thus the forecasts bring only a limited advantage for the daily operations of the salmon farmers or processors. All in all, the salmon spot price remains barely predictable in the short horizons of one to five weeks, and essentially unpredictable beyond the effect of seasonality.

The limited predictability lends support to the hypothesis of weak form efficiency in the salmon spot market. Weak form efficiency results if the salmon spot price is unpredictable using the past prices. The semi-strong form of the efficient market hypothesis implies that price is unpredictable using all publicly available past information, which is also compatible with the findings of this study. Using a number of relevant variables for prediction I was not able to beat the naïve seasonal forecast by a significant margin. Whether this finding is surprising could be debated. On the one hand, the salmon production has become quite concentrated recently; the number of the largest companies that together produce 80% of the Norwegian farmed salmon has shrunk from around 70 in 1997 to 24 in 2013 (Marine Harvest, 2014a). A decrease in the number of market participants may be expected to reduce market efficiency. On the other hand, the salmon spot market has been operating continuously for a few decades now, and the traded volumes are substantial. A number of large, dominating market players have accumulated considerable experience and expertise in this

[12] Except for the naïve no-change forecast that does not predict the direction of price change, and the artificial neural network in the case of one-week-ahead forecasts.

industry. The existence of a futures market for salmon facilitates timely dissemination of price information via the daily price quotes on the futures contracts. Briefly, price predictability may hardly be expected in such a liquid and developed market.

Recall that the spot price definition has changed on 2013 week 14 when the NOS price was replaced by the NASDAQ price; I identify this as a possible source of forecast error. This change illustrates the real-world challenges to forecasting. There are not enough observations in the post-change period to compare the forecasting performance in samples consisting of purely NOS prices to samples made of purely NASDAQ prices; I also find the sample period before the change of the price definition too short to analyse separately. All the multivariate models and especially the models that depend on cross validation would likely perform poorly with training samples shorter than four years (209 weekly observations) as currently used. Therefore, the effect of the NASDAQ price having replaced the NOS price has not been assessed here and remains to be examined in future studies.

The difficulty in predicting the salmon spot price beyond the seasonal fluctuations leaves the high uncertainty in the future spot price of salmon unresolved. For the salmon market participants this means that price hedging remains perhaps the only viable option for managing the price uncertainty. Therefore, analysis of price hedging seems a natural research direction for the future.

6 Acknowledgements

I thank Ole Gjølberg, Gabriel Power, Gunnar Bårdsen, Arnar Mar Buason, Mindaugas Bloznelis, Christian Kascha, members of the Norwegian Center for Commodity Market Analysis, and the community of Cross Validated statistics website for valuable discussions and feedback. Financial support of the Norwegian Research Council is gratefully acknowledged. All remaining errors are mine.

References

ANSCOMBE, F. J. & GLYNN, W. J. 1983. Distribution of the kurtosis statistic b2 for normal samples. *Biometrika,* 70**,** 227-234.

ASCHE, F., GUTTORMSEN, A. G., SEBULONSEN, T. & SISSENER, E. H. 2005. Competition between farmed and wild salmon: the Japanese salmon market. *Agricultural Economics,* 33**,** 333-340.

ASCHE, F., HANSEN, H., TVETERÅS, R. & TVETERÅS, S. 2009. The salmon disease crisis in Chile. *Marine Resource Economics,* 24**,** 405-411.

BAILLIE, R. T., CHUNG, C. F. & TIESLAU, M. A. 1996. Analysing inflation by the fractionally integrated ARFIMA-GARCH model. *Journal of Applied Econometrics,* 11**,** 23-40.

BÅRDSEN, G. & LÜTKEPOHL, H. 2011. Forecasting levels of log variables in vector autoregressions. *International Journal of Forecasting,* 27**,** 1108-1115.

BLIEMEL, F. 1973. Theil's forecast accuracy coefficient: a clarification. *Journal of Marketing Research***,** 10**,** 444-446.

BLOZNELIS, D. 2016. Salmon price volatility: a weight-class-specific multivariate approach. *forthcoming in Aquaculture Economics & Management*.

BURNHAM, K. P. & ANDERSON, D. R. 2004. Multimodel inference: understanding AIC and BIC in model selection. *Sociological Methods & Research,* 33**,** 261-304.

CARREIRA, R. M. 2013. *The causal relationship between salmon prices and share price. Price analysis on the Oslo Stock Exchange.* Masters, The Arctic University of Norway.

CHIGIRA, H. & YAMAMOTO, T. 2009. Forecasting in large cointegrated processes. *Journal of Forecasting,* 28**,** 631-650.

CHRISTOFFERSEN, P. F. & DIEBOLD, F. X. 1998. Cointegration and long-horizon forecasting. *Journal of Business & Economic Statistics,* 16**,** 450-456.

CLEVELAND, R. B., CLEVELAND, W. S., MCRAE, J. E. & TERPENNING, I. 1990. STL: a seasonal-trend decomposition procedure based on loess. *Journal of Official Statistics,* 6**,** 3-73.

D'AGOSTINO, R. B. 1970. Transformation to normality of the null distribution of G1. *Biometrika,* 57**,** 679-681.

DE GOOIJER, J. G. & HYNDMAN, R. J. 2006. 25 years of time series forecasting. *International Journal of Forecasting,* 22**,** 443-473.

DIEBOLD, F. X. 2014. *Forecasting: in economics, business, finance and beyond.* Philadephia, Department of Economics, University of Pennsylvania.

DIEBOLD, F. X. 2015. Comparing predictive accuracy, twenty years later: a personal perspective on the use and abuse of Diebold–Mariano tests. *Journal of Business & Economic Statistics,* 33**,** 1-9.

DIEBOLD, F. X. & MARIANO, R. S. 1995. Comparing predictive accuracy. *Journal of Business & Economic Statistics,* 13**,** 253-263.

DOORNIK, J. A. & OOMS, M. 2004. Inference and forecasting for ARFIMA models with an application to US and UK inflation. *Studies in Nonlinear Dynamics & Econometrics,* 8**,** 1-23.

DUAN, N. 1983. Smearing estimate: a nonparametric retransformation method. *Journal of the American Statistical Association,* 78**,** 605-610.

ELLIOTT, G., ROTHENBERG, T. J. & STOCK, J. H. 1996. Efficient tests for an autoregressive unit root. *Econometrica,* 64**,** 813-836.

ENGLE, R. F. & GRANGER, C. W. 1987. Co-integration and error correction: representation, estimation, and testing. *Econometrica: Journal of the Econometric Society,* 55**,** 251-276.

ENGLE, R. F. & YOO, B. S. 1987. Forecasting and testing in co-integrated systems. *Journal of Econometrics,* 35**,** 143-159.

FISHPOOL. 2014. *NOS and NASDAQ price* [Online]. Fish Pool. Available: http://fishpool.eu/wp-content/uploads/2014/08/FAQ-Nasdaq.pdf [Accessed 2015-03-30].

GHALANOS, A. 2014. rugarch: univariate GARCH models. 1.3-3 ed. Vienna, Austria: R Project for Statistical Computing.

GILBERT, P. D. 1993. *State space and arma models: an overview of the equivalence*, Ottawa, Bank of Canada.

GRANGER, C. W. & JOYEUX, R. 1980. An introduction to long-memory time series models and fractional differencing. *Journal of Time Series Analysis,* 1**,** 15-29.

GU, G. & ANDERSON, J. L. 1995. Deseasonalized state-space time series forecasting with application to the US salmon market. *Marine Resource Economics,* 10**,** 171-185.

GUTTORMSEN, A. G. 1999. Forecasting weekly salmon prices: risk management in fish farming. *Aquaculture Economics & Management,* 3**,** 159-166.

HARVEY, D., LEYBOURNE, S., & NEWBOLD, P. 1997. Testing the equality of prediction mean squared errors. *International Journal of Forecasting,* 13**,** 281-291.

HASTIE, T., TIBSHIRANI, R. & FRIEDMAN, J. 2009. *The elements of statistical learning*, Berlin, Springer.

HYNDMAN, R. J. 2014. Forecasting weekly data. Available from: http://robjhyndman.com/hyndsight/forecasting-weekly-data/ [Accessed 2014-03-05].

HYNDMAN, R. J. 2015. forecast: forecasting functions for time series. 5.8 ed. Vienna, Austria: R Project for Statistical Computing.

HYNDMAN, R. J. & ATHANASOPOULOS, G. 2014. *Forecasting: principles and practice*, OTexts. [Accessed 2015-10-12].

HYNDMAN, R. J. & KOEHLER, A. B. 2006. Another look at measures of forecast accuracy. *International Journal of Forecasting,* 22**,** 679-688.

HOERL, A. E. & KENNARD, R. W. 1970. Ridge regression: biased estimation for nonorthogonal problems. *Technometrics,* 12**,** 55-67.

HOFFMAN, D. L. & RASCHE, R. H. 1996. Assessing forecast performance in a cointegrated system. *Journal of Applied Econometrics,* 11**,** 495-517.

JOHANSEN, S. 1991. Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models. *Econometrica: Journal of the Econometric Society,* 59**,** 1551-1580.

KONISHI, S. & KITAGAWA, G. 2008. *Information criteria and statistical modeling*, Berlin, Springer.

KOOPMAN, S. J., OOMS, M. & CARNERO, M. A. 2007. Periodic seasonal Reg-ARFIMA–GARCH models for daily electricity spot prices. *Journal of the American Statistical Association,* 102**,** 16-27.

LIN, B. H., HERRMANN, M., LIN, T. Y. & MITTELHAMMER, R. C. 1989. Forecasting the price of farmed Atlantic salmon: an integrated econometric and time series approach. *Agribusiness,* 5**,** 477-488.

LJUNG, G. M. & BOX, G. E. 1978. On a measure of lack of fit in time series models. *Biometrika,* 65**,** 297-303.

LÖF, M. & FRANSES, P. H. 2001. On forecasting cointegrated seasonal time series. *International Journal of Forecasting,* 17**,** 607-621.

LUMLEY, T. & MILLER, A. 2009. leaps: regression subset selection. 2.9 ed. Vienna, Austria: R Project for Statistical Computing.

LÜTKEPOHL, H. 2007. *New introduction to multiple time series analysis*, Berlin, Springer.

MADDALA, G. S. & KIM, I.-M. 1998. *Unit roots, cointegration, and structural change*, Cambridge, Cambridge University Press.

MAKRIDAKIS, S., ANDERSEN, A., CARBONE, R., FILDES, R., HIBON, M., LEWANDOWSKI, R., NEWTON, J., PARZEN, E. & WINKLER, R. 1982. The accuracy of extrapolation (time series) methods: results of a forecasting competition. *Journal of Forecasting,* 1**,** 111-153.

MAKRIDAKIS, S. & HIBON, M. 1995. Evaluating accuracy (or error) measures. *INSEAD Working Papers.* Fontainebleau: INSEAD.

MARCELLINO, M., STOCK, J. H. & WATSON, M. W. 2006. A comparison of direct and iterated multistep AR methods for forecasting macroeconomic time series. *Journal of Econometrics,* 135**,** 499-526.

MARINE HARVEST 2014a. Salmon farming industry handbook 2014. Bergen: Marine Harvest.

MARINE HARVEST. Scenarios for global salmon production next five years. Havbrukskonferansen, 2014-11-26 2014b Oslo.

MEADE, N. & SMITH, I. M. 1985. ARARMA vs ARIMA—a study of the benefits of a new approach to forecasting. *Omega,* 13**,** 519-534.

MINCER, J. A. & ZARNOWITZ, V. 1969. The evaluation of economic forecasts. *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance.* Cambridge: National Bureau of Economic Research.

NEWBOLD, P. & HARVEY, D. I. *In:* CLEMENTS, M. P. & HENDRY, D. F. 2008. (eds.) *A companion to economic forecasting*, Hoboken, John Wiley & Sons.

NEWEY, W. K. & WEST, K. D. 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55**,** 703-708.

OSLO BØRS. *OSLO Seafood Index* [Online]. Available: http://www.oslobors.no/ob_eng/markedsaktivitet/#/details/OSLSFX.OSE/overview [Accessed 2015-03-25].

PALL, M., NORBERG, B., ANDERSON, E. & TARANGER, G. L. 2006. Kjønnsmodning hos atlantisk laks (Salmo salar). *In:* SVÅSAND, T., BOXASPEN, K., DAHL, E. & JØRGENSEN, L. L. (eds.) *Fisken og havet.* Bergen: Havforskningsinstitutet.

PARZEN, E. 1982. ARARMA models for time series analysis and forecasting. *Journal of Forecasting,* 1**,** 67-82.

PESARAN, M. H., PICK, A. & TIMMERMANN, A. 2011. Variable selection, estimation and inference for multi-period forecasting problems. *Journal of Econometrics,* 164**,** 173-187.

SAID, S. E. & DICKEY, D. A. 1984. Testing for unit roots in autoregressive-moving average models of unknown order. *Biometrika,* 71**,** 599-607.

SHAPIRO, S. S. & FRANCIA, R. 1972. An approximate analysis of variance test for normality. *Journal of the American Statistical Association,* 67**,** 215-216.

SMITH, J. & WALLIS, K. F. 2009. A simple explanation of the forecast combination puzzle. *Oxford Bulletin of Economics and Statistics,* 71**,** 331-355.

STONE, M. 1977. An asymptotic equivalence of choice of model by cross-validation and Akaike's criterion. *Journal of the Royal Statistical Society. Series B (Methodological),* 39**,** 44-47.

STRAUME, H.-M. 2014. Currency invoicing in Norwegian salmon export. *Marine Resource Economics,* 29**,** 391-409.

THE FISH SITE. 2015. *Record Salmon Exports for Norway in 2014* [Online]. The Fish Site. Available: http://www.thefishsite.com/fishnews/24886/record-salmon-exports-for-norway-in-2014/ [Accessed 2015-03-30].

TIBSHIRANI, R. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological),* 58**,** 267-288.

TIMMERMANN, A. 2006. Forecast combinations. *In:* ELLIOTT, G., GRANGER, C. W. J. & TIMMERMANN, A. (eds.) *Handbook of Economic Forecasting.* 1 ed.

VENABLES, W. N. & RIPLEY, B. D. 2002. *Modern applied statistics with S*, Berlin, Springer.

VUKINA, T. & ANDERSON, J. L. 1994. Price forecasting with state-space models of nonstationary time series: case of the Japanese salmon market. *Computers & Mathematics with Applications,* 27**,** 45-62.

ZOU, H. & HASTIE, T. 2005. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology),* 67**,** 301-320.

Paper 3

Hedging salmon price risk

Daumantas Bloznelis[1]

School of Economics and Business

Norwegian University of Life Sciences

Abstract

Salmon price is highly volatile and difficult to predict, which obscures planning decisions and raises financing costs for the salmon market participants. I consider hedging the uncertainty in the future spot price of salmon using salmon futures, live cattle, soybean meal and soybean oil futures, and shares of salmon farming companies listed on the Oslo Stock Exchange. I propose a new objective function for hedging and a new measure of hedging effectiveness tailored to partly predictable spot prices of perishable commodities. Among all the hedging instruments and their combinations, the salmon futures contract offers the highest hedging effectiveness; however, low liquidity may limit its applicability in practice.

Keywords: salmon price, price risk, hedging, hedging effectiveness, optimal hedge ratio, futures

[1] Mailing address: School of Economics and Business, Norwegian University of Life Sciences, P.O. Box 5003, NO-1432 Ås, Norway. E-mail: daumantas.bloznelis@nmbu.no.

1 Introduction

The spot price of Norwegian farmed Atlantic salmon is highly volatile as compared with other soft commodities (Bloznelis, 2016a); it is also difficult to predict (Bloznelis, 2016b). Large unexpected price changes create difficulties in business management for salmon farmers, processors, exporters, retailers and other market participants. They also and make financing expensive as high uncertainty implies high required return via the capital asset pricing model (Sharpe, 1964, Lintner, 1965). The high uncertainty is perceived negatively by the market participants (Jensen, 2013) as it requires them to be flexible and financially solid enough to withstand large unexpected increases in costs or decreases in revenues due to changes in the spot price of salmon.

There are four ways to mitigate the effect of the high price uncertainty. First, the original asset's price may be fixed by using forward contracts; bilateral contracts allow buying or selling salmon in advance of the actual delivery. However, delayed delivery brings in a new source of risk; one of the contracting parties may not be able to deliver on its obligations when the settlement time comes (using third-party clearing service could prevent the associated losses for an extra cost). Also, it need not always be easy to find counterparty as forward contracts are specific rather than standardized. Second, an alternative to forward contracts is futures contracts. Counterparty risk is eliminated by clearing and finding counterparty is typically easier since the contracts are standardized. However, the salmon futures market is known to suffer from low liquidity. Third, cross-hedging can be considered an option. It amounts to using other financial instruments such as futures contracts on other assets as hedging tools. In the case of salmon, futures on substitute products such as beef or input products such as soybeans could be considered. Also, shares of salmon farming companies listed on the Oslo Stock Exchange may provide a hedge for the salmon price. The revenues of a salmon farming company depend to a large extent on the salmon price; therefore, the share price may be correlated with the current and the expected spot price of salmon and thus constitute an effective hedge. Fourth, accurate price forecasts would provide the market players with time to adapt to the upcoming price changes. However, obtaining accurate salmon price forecasts is challenging, if not impossible (Bloznelis, 2016b).

Hedging salmon price risk has received little attention in the academic literature. Bergfjord (2007) discusses hedging by salmon futures in the context of the perspectives for a futures market for salmon; it is a hypothetical discussion with no empirical results. Bergfjord (2009) reports a survey on risk perception and risk management of Norwegian salmon farmers conducted in 2005. Salmon farmers appear to be only moderately risk averse, but they consider the future price of salmon to be the most important source of risk.

In this article I define an objective function for hedging the salmon price risk and propose a new measure of hedging effectiveness relevant for commodities with partly predictable prices. I then suggest instruments and describe their use for hedging the uncertainty in the future spot price of salmon four, eight and 13 weeks ahead. Finally, I derive optimal hedge ratios and examine out-of-sample hedging performance.

The remainder of the article is structured as follows. Chapter 2 provides an introduction to the salmon market. Chapter 3 presents hedging in general and considers hedging the future spot price of salmon in particular. Methodology of the paper is laid out in Chapter 4, succeeded by a review of the data in Chapter 5. Empirical results and a discussion conclude.

2 Salmon market

Commercial farming of salmon started in 1970s in Norway.[2] The industry expanded rapidly so that by late 1990s the aquaculture production volume had soared past the wild catch volume; see Figure 1[3]. In 2014, the global farmed salmon production was at two million metric tons valued roughly at nine billion Euros. About 55% of the volume is of Norwegian origin; the other major salmon producers are Chile, Canada and Scotland. Farmed salmon is consumed all over the world; the largest markets are the EU (51% of global consumption) and the U.S. (24%).

Figure 1 Historical farmed salmon production volumes, Norwegian (dark fill) and global, 1980-2013 (million tonnes)

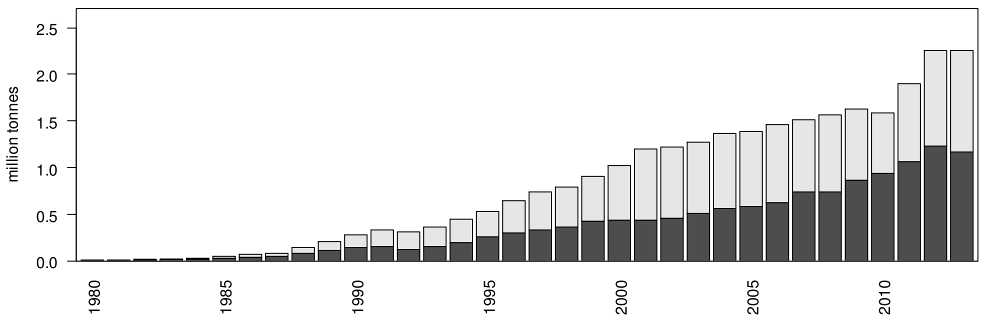

[2] This section is largely based on the information in Marine Harvest (2015).

[3] The data is taken from Food and Agriculture Organization of the United Nations and can be accessed at http://www.fao.org/figis/servlet/SQServlet?file=/work/FIGIS/prod/webapps/figis/temp/hqp_5437944154593549544.xml&outtype=html.

The price of the Norwegian farmed Atlantic salmon experienced a continuous decrease from the 1980s until the break of the millennia; see Figure 2[4]. The decline is attributed to technological improvement and fast supply growth (Asche and Bjorndal, 2011, p. 43-48). The price trend reversed in the early 2000s likely because of slower technological progress, reduced supply growth due to limited availability of suitable production sites, increasing raw material costs (Oglend and Asche, 2015) and increased demand from the emerging markets.

Figure 2 Norwegian farmed Atlantic salmon price, 1980-2013 (NOK/kg)

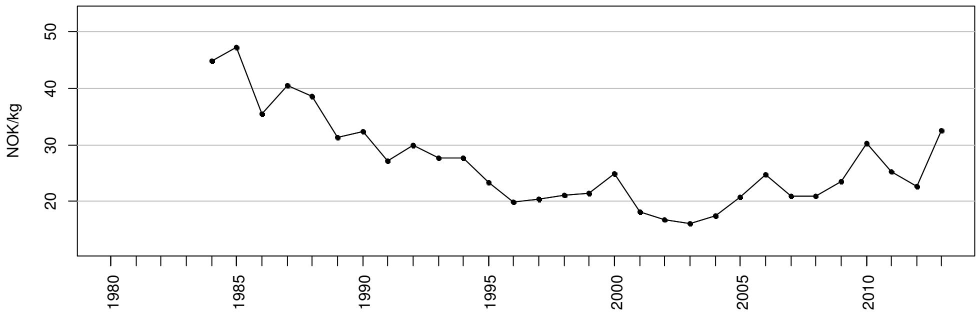

Price uncertainty may arise due to supply and demand factors. On the supply side there are several possible production risks, e.g. infectious disease (an outbreak of infectious salmon anemia decimated the Chilean salmon production in late 2000s), biological conditions (such as presence or absence of toxic blooming algae) and the uncertainty in future water temperature (suboptimal temperature leads to slower growth). On the demand side, price uncertainty may be affected by unexpected changes in consumer tastes resulting from positive or negative publicity, e.g. reports on health gains from a salmon-rich diet or concerns over pollution from the salmon farms; or changes in price of substitutes or complements, or purchasing power. I conjecture that supply factors dominate the demand factors in causing large unexpected price fluctuations in the short term. However, a formal treatment would be needed to confirm the conjecture.

A substantial part of salmon is sold on the spot market, but forward contracting is also substantial (Larsen and Asche, 2011). Since mid-2006 there exists a futures exchange for salmon, Fish Pool, based in Bergen, Norway. Monthly contracts for one through 60 months ahead are available for trading. The contracts are cash settled; thus, no physical fish is bought or sold at Fish Pool. The contracts' reference price is the "Fish Pool Index" (FPI; since 2015 also denoted as FPSA15), a

[4] The data is taken from Food and Agriculture Organization of the United Nations at http://www.fao.org/figis/servlet/SQServlet?file=/work/FIGIS/prod/webapps/figis/temp/hqp_5437944154593549544.xml&outtype=html and from Norges Bank at http://www.norges-bank.no/Statistikk/Valutakurser/valuta/USD/.

weighted average of three (previously four) prices that may be considered reasonably representative of the spot price of 3-6 kg salmon. The futures market suffers from low liquidity as its turnover matches only about a tenth of the physical market volume (Fish Pool, 2015) and the trades are rather infrequent (as evidenced by confidential data obtained from Fish Pool). Fish Pool also provides clearing services for salmon forward contracts, and it is not inconceivable that futures prices are used as settlement prices in some of the forward contracts.

3 Hedging

3.1 The idea of hedging

Hedging is "making an investment to reduce the risk of adverse price movements in an asset" (Investopedia). The exact purpose of hedging may be defined by specifying a loss function in terms of the anticipated outcomes resulting from the hedged position. When hedging the price uncertainty of an asset h periods ahead using a hedging instrument, a popular choice of the loss function is the expected (as of time t) squared portfolio return

$$E_t\left(\left((y_{t+h} - \beta x_{t+h}) - (y_t - \beta x_t)\right)^2\right) \tag{1a}$$

where y is the price of the given asset, x is the price of the hedging instrument, β is the portfolio weight of the hedging instrument, also known as the hedge ratio, and $y - \beta x$ is the price of the hedge portfolio. The hedge ratio, β, gives the exposure to the hedging instrument's price as a fraction (or a multiple) of the exposure to the price of the original asset. Given a loss function in equation (1a), the goal of the hedger is to minimize it with respect to the hedge ratio, β:

$$E_t\left(\left((y_{t+h} - \beta x_{t+h}) - (y_t - \beta x_t)\right)^2\right) \to \min_{\beta} \tag{1b}$$

The ratio β^* that minimizes the loss function is called the optimal hedge ratio (OHR) or the risk-minimizing hedge ratio. To be able to solve for β^*, a model for the expected prices y and x is needed. The hedging problem can also be formulated in terms of an objective function in place of a loss function. Chen et al. (2003) provide a review of solving for and estimating optimal hedge ratios under different assumptions and objective functions. Lence (1995) and Lence (1996) cited in Garcia and Leuthold (2004) show that in practice (under transaction costs, production uncertainty etc.) optimal hedge ratio may as well be zero.

The problem of a salmon farmer is, how to reduce the uncertainty around the future spot price (or equivalently, the return on the spot price) of salmon. If hedging is to be used, the problem can be split into two stages. First, finding an effective hedging instrument. Second, estimating the hedge ratio β that would minimize the uncertainty around the future value of the hedge portfolio. I will measure the uncertainty as the expected (as of time t) squared forecast error. Formally, the farmer's objective function is

$$E_t\left(\left((s_{t+h}-\beta f_{t+h})-E_t(s_{t+h}-\beta f_{t+h})\right)^2\right)\rightarrow \min_{\beta} \quad (2)$$

where s is the spot price of salmon and f is the price of the hedging instrument (f stands for "futures" and will be the natural notation for all hedging instruments except for the share prices). Alternatively, the problem could be formulated from a buyer's perspective. Salmon processor, exporter or retailer buying salmon directly from the salmon farmer and paying the spot price also share the objective function given in (2).

3.2 Relevant hedging instruments

Let us first review the hedging instruments that could be relevant for the spot price of salmon and then explore the estimation of the optimal hedge ratio, β. I will consider the following hedging instruments: (1) salmon futures; (2) shares of salmon farming companies; (3) live cattle, soybean meal and soybean oil futures. Salmon futures is a natural hedging instrument for the spot price as the underlying of a futures contract is the average price of (one tonne of) salmon over a given month, as measured by FPI. The FPI roughly equals the spot price; their correlation is above 0.98 both in levels and first differences for 2007-2015. Hence, a salmon futures contract constitutes a nearly perfect hedging instrument for a farmer selling an equal amount of fish every week of the month. The hedging effectiveness may drop due to deviations in volume sold from week to week, variations in fish quality resulting in discrepancies between the FPI and the price obtained by the farmer, etc.

Suppose the fish is to be sold on the spot market on the first week of January. I will examine three alternative uses of the salmon futures contracts to hedge this physical position. First, a January contract can be sold in advance and bought back on the first week of January; I will briefly denote this strategy F0. Second, a January contract can be sold in advance and kept until expiration (maturity), which is at the end of January; this alternative will be denoted M (due to "maturity"). Third, a February contract can be sold in advance and bought back on the first week of January (F1). The first strategy may provide an effective hedge as the price of a futures contract close to maturity should be highly correlated to the seasonally-adjusted spot price. However, the strategy may be difficult to implement due to the lack of liquidity in the futures market, especially since the trading volume tends to decrease when the contract approaches the expiration date. Also, prices of contracts close to maturity may be volatile and thus less correlated with the spot price making the contracts less effective as hedging instruments. The other two strategies avoid this issue. The second alternative, M, does not suffer from the liquidity problem as the contract is not bought back but rather kept until expiration. It also is cost effective since commissions associated with buying the contract back are avoided. However, the price at the expiration date (which equals the average FPI over a whole month, in this case, January) may be less correlated with the spot price of the particular week when the physical fish is sold (the first week of January). The third alternative, F1, also circumvents the expiration problem at the expense of potentially lower price correlation; using F1 is common practice in some futures markets (Hull, 2012, p. 55). It is difficult to foresee which of the three strategies should be the most effective in the case of salmon; their relative performance will be revealed by empirical analysis in Chapter 6.

Hedging the price uncertainty for a particular week using a monthly futures contract is of course less effective than hedging the price uncertainty for a whole month, as long as the production volume is approximately equally distributed across the different weeks of the month. This is because the underlying of a futures contract closely matches the physical position for a whole month rather than a particular week. If the price faced by a market participant closely matches the market average price, the salmon futures contract kept until maturity would offer a nearly perfect hedge for a whole month's physical position. More formally, we would have $s_{t+h} \approx f_{t+h}$ and $E_t(s_{t+h}) \approx E_t(f_{t+h})$, and hence $s_{t+h} - \beta f_{t+h} \approx 0$ and $E_t(s_{t+h} - \beta f_{t+h}) \approx 0$ for $\beta = 1$. Thus there are good grounds to expect $E_t\left(\left((s_{t+h} - \beta f_{t+h}) - E_t(s_{t+h} - \beta f_{t+h})\right)^2\right) \approx 0$ which means the objective function given in equation (2) would essentially be optimized. However, if the production volume is distributed unequally across the weeks in a month, the physical position will not match closely the underlying of the futures contract; hence, hedging with salmon futures will likely be less effective.

Now consider the price of a salmon farming company's share as a hedging instrument. The share price should naturally depend on the spot price of salmon. Ceteris paribus, the higher the spot price, the higher the income and the profit of the company, the more valuable the company's share; see e.g. Carreira (2013) for an empirical study of the relation between the salmon price and the share prices of salmon farming companies. While most of the salmon farming companies are not traded intensively and have medium to low liquidity, Marine Harvest ASA stands out as one of the most liquid assets on the Oslo Stock Exchange; it has long been included in the OBX, an index of the most liquid companies. The share price of Marine Harvest is thus a viable hedging instrument. Note that buying a share and selling it later is simpler and cheaper than short-selling the share and buying it back later; therefore, this hedging instrument is more readily suited for buyers rather than sellers of physical salmon. Also, share prices of salmon farming companies are driven by many factors aside from the spot price of salmon. The effect of market sentiment, interest rates etc. may reduce the correlation between the spot price of salmon and the share price of Marine Harvest. Therefore, one may adjust the share price by subtracting a scaled market index expecting the adjusted share price to be more highly correlated with the spot price of salmon. The adjusted share price is a technically feasible hedging instrument as there exists an exchange traded fund following the Oslo Stock Exchange Benchmark Index (OSEBX), and there also exist futures contracts on the index; combining the two yields the adjusted share price. To find the relevant scale of the market index one may run a regression of the share price on the stock index and extract the slope coefficient. The adjusted share price is obtained by subtracting the stock index times the slope coefficient from the share price. I will explore the hedging performance of both the raw share price of Marine Harvest and the OSEBX-adjusted share price.

Beef may be considered a substitute for salmon. Hence, we may expect the price of salmon to move in the same direction as the price of beef. A futures contract tied to beef price could thus be a relevant hedging instrument. There are two futures contracts related to beef, for feeder cattle and for live cattle. The former is written on the price of young calves that are not yet ready for slaughtering and thus is of similar relevance. Meanwhile, the latter is tied to the price of grown-up animals ready to be slaughtered and converted into beef and hence is the relevant contract. The live cattle contract is actively traded on the Chicago Mercantile Exchange. There are six monthly contracts per year maturing in February, April, June, August, October and December. When hedging a physical position in a month without a live cattle futures contract, the nearest subsequent contract will be used.

Feed cost is the largest variable cost in modern salmon farming (Marine Harvest, 2015). With feed cost increasing or decreasing, the salmon price may be expected to follow suit. Unsurprisingly, the prices of fishmeal and soybeans have been found to be cointegrated with the price of salmon (Oglend and Asche, 2015). This suggests fishmeal and soybean futures could be used as hedging instruments. However, a fishmeal futures contract does not exist; hence, it is not a feasible hedging instrument. I will only investigate soybean meal and soybean oil futures traded at Chicago Board of Trade (CBOT). The monthly contracts mature eight times a year: in January, May, May, July, August, September, October and December.

Both for the live cattle and for the soybean futures, we may expect lower hedging effectiveness for the months without contracts (six for live cattle and four for soybean meal and oil), and hence overall. Another nuisance that may lower the hedging effectiveness is that these futures are denominated in U.S. dollars. Hence, the relation between the salmon price and the futures prices is subject to exchange rate fluctuations, and the procedure of hedging is more involved and costly due to the need to exchange currencies.

Alternative hedging instruments could be, for example, options on the FPI and options on the share price of Marine Harvest. The former suffer from low liquidity even more than the salmon futures do, and hence do not seem to be a feasible hedging instrument. Meanwhile, the latter are quite liquid and could in principle be considered. However, hedging with options may be more expensive than hedging with futures and is not explored further in this study.

The salmon price considered so far was the spot price averaged over the different sizes of fish. Seven weight classes of salmon are distinguished in the historical price data: 1-2 kg, 2-3 kg, …, 6-7 kg and 7+ kg. Due to seasonality in both the supply and the demand for salmon, the relative production volumes of the different weight classes vary in one-year cycles; the relative prices do as well (Asche and Guttormsen, 2001, Marine Harvest, 2015). Thus effectively there are seven different objects to hedge. However, hedging a particular weight class of salmon is not generally relevant as the production schedule is relatively similar across the major salmon producers (Guttormsen, 2015) so that all of them produce salmon of different weight classes in similar proportions. Nevertheless, there may be cases when it is known in advance that a farmer will have to sell an amount of salmon that has a different weight class mix than is typical for the season. For instance, this may happen if the amount of salmon on a farming site exceeds the allowable maximum, and keeping the salmon unslaughtered would incur a fine from the regulatory agencies. Then a need to hedge particular weight classes arises. Therefore, I will examine weight-class-specific hedging as well as hedging the average price.

4 Methods

To be able to estimate the optimal hedge ratio as per equation (2), given a spot position to be hedged and a hedging instrument, first we need to model the joint behaviour of the spot price and the instrument's price. A predictive model for the conditional mean vector and the conditional variance matrix of the bivariate price system will provide a forecast of the conditional mean and variance, which will then allow estimating the optimal hedge ratio. I will use rolling windows within the original sample and will model the data and derive the optimal hedge ratio separately in each window. The hedge ratio will be applied to the data immediately following the rolling window to construct the hedge portfolio and assess its out-of-sample hedging performance.

4.1 Modelling and forecasting price

Salmon price is known to be seasonal (Asche and Guttormsen, 2001); recent evidence that seasonality is a genuine feature rather than a spurious pattern can be found in Bloznelis (2016b) where the seasonal components extracted in sample are successfully used in forecasting the spot price of salmon out of sample. I use a seasonal adjustment procedure following Hyndman (2014) as described in Bloznelis (2016b). I also seasonally adjust the salmon, live cattle, soybean meal and soybean oil futures prices using the same method.

The small size of the rolling windows limits the choice of models. Sophisticated forecasting models are not likely to fare well when estimated in small samples because of their high variance. Therefore, I choose two basic conditional mean models. First, both prices are assumed to have expected changes of zero:

$$\begin{aligned} s_t &= s_{t-1} + u_t \\ f_t &= f_{t-1} + v_t; \end{aligned} \qquad (3)$$

here s_t denotes the spot price, f_t denotes the hedging instrument's price, and u_t and v_t denote zero-mean, non-autocorrelated error terms. Second, the prices are assumed to follow a vector error correction model (VEC model, or VECM). If the spot price and the hedging instrument's price are assumed to be cointegrated, a VECM is the appropriate model for the system:

$$\Delta \begin{pmatrix} s_t \\ f_t \end{pmatrix} = \begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix} (1 \quad \beta) \begin{pmatrix} s_{t-1} \\ f_{t-1} \end{pmatrix} + \Gamma_1 \Delta \begin{pmatrix} s_{t-1} \\ f_{t-1} \end{pmatrix} + \cdots + \Gamma_{\mathrm{p}} \Delta \begin{pmatrix} s_{t-p} \\ f_{t-p} \end{pmatrix} + \begin{pmatrix} u_t \\ v_t \end{pmatrix}; \qquad (4)$$

here Δ denotes difference operator such that $\Delta x_t = x_t - x_{t-1}$; α_1, α_2 are loading coefficients; $(1 \quad \beta)$ forms the (transposed) cointegrating vector such that for a pair of integrated processes (s_t, f_t), $(1 \quad \beta)\begin{pmatrix} s_t \\ f_t \end{pmatrix}$ is a stationary process; Γ_1 through Γ_{p} are 2×2 coefficient matrices; and u_t and v_t are zero-mean, non-autocorrelated error terms. When salmon futures will be used as a hedging instrument, I will not estimate the cointegration rank and the cointegrating vectors in each rolling window but rather arbitrarily impose presence of cointegration with a (transposed) cointegrating vector $(1 \quad \beta) = (1, \quad -1)$. This is to prevent the possible effect of imprecise estimation of the cointegrating vector due to the small sample size. (I also tried an unrestricted VECM, and the modelling results were very similar to the case of the arbitrarily imposed cointegration vector.) The lag order or the VECM is selected using Akaike's information criterion (AIC). AIC-based choice tends to yield the model that produces the smallest squared one-step-ahead forecast error among the set of candidate models (Konishi and Kitagawa, 2008, p. 249-250), which is relevant when selecting a model for forecasting. The conditional mean models will produce point forecasts and also in-sample residuals. The latter will be used for modelling the conditional variance.

Once the residuals from the conditional mean model are obtained, the residual variance matrix can be formed. The variance matrix might be either constant or time varying. If it is assumed to be constant, the sample variance matrix will be used as its estimate. If it is time varying, some version of multivariate generalized autoregressive conditional heteroskedasticity (GARCH; Bollerslev (1986)) model could be employed to describe the development of the individual variances and the covariance over time. In Bloznelis (2016a), the time-varying variance matrices of the weight-class-specific spot prices were modeled using univariate GARCH models combined with a dynamic conditional correlation (DCC) model (Engle, 2002). This is a popular approach in financial applications due to the model's relatively simple specification and fast estimation; I will use it here as well.

Let u_t denote the error term (the innovation) of the conditional mean model and let σ_t^2 denote the conditional variance of u_t given the information up until time period $t-1$. A GARCH(1,1) model for the conditional variance of u_t is specified as follows:

$$\begin{aligned} u_t &= \sigma_t \varepsilon_t \\ \sigma_t^2 &= \omega + \alpha u_{t-1}^2 + \beta \sigma_{t-1}^2; \end{aligned} \tag{5}$$

here ε_t is an $i.i.d.(0,1)$ random variable; ω is a basis level of the conditional variance, σ_t^2, such that the conditional variance never drops below ω. This model implies the conditional variance is deterministic rather than stochastic; it is completely determined by the last period's conditional variance and the last innovation from the conditional mean model. The index (1,1) in the model name suggests that lagged innovation and lagged conditional variance from one period before are used. In principle, higher-order memory indexed as (q,r) could be used; however, the (1,1) specification is often found to be sufficient (Hansen and Lunde, 2005) and is the most widely used in practice.

A bivariate DCC(1,1) model treats the conditional correlation between a pair of innovations in an analogous way as to how a GARCH(1,1) model specifies the conditional variance of an innovation. Given the standardized innovations $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$ from the GARCH models for the spot price and the hedging instrument's price, respectively, define conditional quasi-correlation at time t as

$$q_{1,2,t} = (1-\gamma-\delta)\rho + \gamma \varepsilon_{1,t-1}\varepsilon_{2,t-1} + \delta q_{1,2,t-1} \tag{6}$$

with unconditional correlation ρ and lagged conditional quasi-correlations q_{t-1}. Then the proper conditional correlation is obtained as scaled conditional quasi-correlation,

$$\rho_{1,2,t} = \frac{q_{1,2,t}}{\sqrt{q_{1,1,t}}\sqrt{q_{2,2,t}}}. \tag{7}$$

(The notation $\rho_{1,2,t}$ could be simplified to ρ_t as the conditional variance matrix of a bivariate system only has one unique off-diagonal element; however, the explicit subscript is kept for completeness.) Scaling ensures the conditional correlation lies strictly between negative one and one. Normally the scaling should have a negligible effect, so that for illustrative purposes one could think of the conditional correlation as approximately following a GARCH-type of model

$$\rho_{1,2,t} \approx (1-\gamma-\delta)\rho + \gamma \varepsilon_{1,t-1}\varepsilon_{2,t-1} + \delta \rho_{1,2,t-1}. \tag{8}$$

Similarly to GARCH, the index (1,1) for DCC indicates use of only one period's lag of the standardized innovations and the conditional (quasi-) correlation. While higher lag orders could be considered, the small size of the rolling window prompts the use of the most parsimonious (1,1) specification.

I estimate the conditional mean and variance models in two stages. This approach is less efficient than simultaneous estimation in the case of time-varying conditional variance matrices, but it is simpler to implement and computationally less demanding.

Combining the two alternatives for the conditional mean specification and the two alternatives for the conditional variance, four models are obtained. Evaluating their performance out of sample will show which of them is to be preferred. However, in practice the model has to be chosen before the outcome is observed. Thus in a real-world application the strategy could be to choose the model having the best historical performance or the one with the lowest AIC value; for the latter option to be valid, the conditional mean and variance models would need to be estimated simultaneously.

4.2 Optimal hedge ratio

A natural candidate for the optimal hedge ratio due to equation (2) for hedging h weeks ahead, β_h^*, is

$$\beta_h^* = \frac{Cov_t(s_{t+h} - s_t, f_{t+h} - f_t)}{Var_t(f_{t+h} - f_t)} = \frac{Cov_t(\sum_{i=1}^h u_{t+i}, \sum_{i=1}^h v_{t+i})}{Var_t(\sum_{i=1}^h v_{t+i})} \tag{9}$$

where (as before) u_t is the innovation from the conditional mean model for s_t; v_t is the innovation from the conditional mean model for f_t; and subscript t to covariance and variance operators indicates conditioning on information available up to time t. Assuming that the innovations from different time periods are conditionally uncorrelated, $Cov_t(u_t, u_s) = Cov_t(u_t, v_s) = Cov_t(v_t, v_s) = 0$ for $t \neq s$, we obtain

$$\beta_h^{**} = \frac{\sum_{i=1}^h Cov_t(u_{t+i}, v_{t+i})}{\sum_{i=1}^h Var_t(v_{t+i})}. \tag{10}$$

To make the expression operational, population quantities will be substituted with their sample counterparts. In the case of constant conditional variance, the sample error variance matrix will provide both quantities. That is, $Cov_t(u_{t+i}, v_{t+i})$ will be estimated by the sample covariance between u and v, and $Var_t(v_{t+i})$ will be estimated by the sample variance of v

$$\begin{aligned}\hat{\beta}_h^{const.var.} &= \frac{\sum_{i=1}^h \widehat{Cov}_t(u, v)}{\sum_{i=1}^h \widehat{Var}_t(v)} = \frac{\sum_{i=1}^h \left(\frac{1}{T-1}\sum_{t=1}^T \left(u_t - \frac{1}{T}\sum_{t=1}^T u_t\right)\left(v_t - \frac{1}{T}\sum_{t=1}^T v_t\right)\right)}{\sum_{i=1}^h \left(\frac{1}{T-1}\sum_{t=1}^T \left(v_t - \frac{1}{T}\sum_{t=1}^T v_t\right)^2\right)} \\ &= \frac{\sum_{i=1}^h \left(\frac{1}{T-1}\sum_{t=1}^T u_t v_t\right)}{\sum_{i=1}^h \left(\frac{1}{T-1}\sum_{t=1}^T v_t^2\right)}\end{aligned} \tag{11}$$

where for the third equality I use the property of the ordinary least squares estimator that the sample means of u and v are zero. In the case of time-varying conditional variance, forecasts of $Cov_t(u_{t+i}, v_{t+i})$ and $Var_t(v_{t+i})$ from the DCC-GARCH model will be substituted into equation (10) to obtain

$$\hat{\beta}_h^{DCC-GARCH} = \frac{\sum_{i=1}^h \widehat{Cov}_t(u_{t+i}, v_{t+i})}{\sum_{i=1}^h \widehat{Var}_t(v_{t+i})}. \tag{12}$$

where hats denote the predicted values from the DCC-GARCH model.

4.3 Measuring hedging effectiveness

Given the conditional mean and variance predictions for the bivariate price systems and the estimated optimal hedge ratios, hedge portfolios can be formed and their returns predicted. Comparing these with the realized out-of-sample returns will allow assessing the hedging effectiveness. In order to match the objective function for hedging given by equation (12) with a measure of hedging effectiveness, let me define hedging effectiveness as the relative reduction in the return uncertainty due to replacing the spot position with the hedge portfolio. That is, hedging effectiveness compares the expected mean squared forecast error of the portfolio return with the expected mean squared forecast error of the spot return in the following way:

$$RRMSFE = \frac{E_t\left(\left(s_{t+h} - E_t(s_{t+h})\right)^2\right) - E_t\left(\left((s_{t+h} - \beta f_{t+h}) - E_t(s_{t+h} - \beta f_{t+h})\right)^2\right)}{E_t\left(\left(s_{t+h} - E_t(s_{t+h})\right)^2\right)}. \quad (13)$$

$RRMSFE$ stands for "relative reduction in mean squared forecast error". The expression in (13) gives a population measure, while its sample counterpart is

$$\begin{aligned}
&\widehat{RRMSFE} \\
&= \frac{\frac{1}{W}\sum_{w=1}^{W}\left(s_{t+h,w} - \hat{s}_{t+h,w}\right)^2 - \frac{1}{W}\sum_{w=1}^{W}\left(\left(s_{t+h,w} - \hat{\beta}_h f_{t+h,w}\right) - \left(\hat{s}_{t+h,w} - \hat{\beta}_h \hat{f}_{t+h,w}\right)\right)^2}{\frac{1}{W}\sum_{w=1}^{W}\left(s_{t+h,w} - \hat{s}_{t+h,w}\right)^2} \quad (14) \\
&= \frac{\sum_{w=1}^{W}\left(s_{t+h,w} - \hat{s}_{t+h,w}\right)^2 - \sum_{w=1}^{W}\left(\left(s_{t+h,w} - \hat{\beta}_h f_{t+h,w}\right) - \left(\hat{s}_{t+h,w} - \hat{\beta}_h \hat{f}_{t+h,w}\right)\right)^2}{\sum_{w=1}^{W}\left(s_{t+h,w} - \hat{s}_{t+h,w}\right)^2}
\end{aligned}$$

where the subscript w indexes the rolling windows, $\hat{s}_{t+h,w}$ stands for an h-step ahead conditional mean forecast of s_{t+h} made at time t in window w, and $\hat{f}_{t+h,w}$ stands for the analogous conditional mean forecast of f_{t+h}.

$RRMSFE$ lies in the interval $(-\infty, 1]$. Larger values indicate greater hedging effectiveness. Negative values signal that hedging is detrimental, i.e. the uncertainty associated with the portfolio return is greater than the uncertainty over the spot return, $E_t\left(\left((s_{t+h} - \beta f_{t+h}) - E_t(s_{t+h} - \beta f_{t+h})\right)^2\right) > E_t\left(\left(s_{t+h} - E_t(s_{t+h})\right)^2\right)$. Positive values indicate that hedging helps reduce the uncertainty, $E_t\left(\left((s_{t+h} - \beta f_{t+h}) - E_t(s_{t+h} - \beta f_{t+h})\right)^2\right) < E_t\left(\left(s_{t+h} - E_t(s_{t+h})\right)^2\right)$. Ideally, the uncertainty over the portfolio return would be zero, and then $RRMSFE$ reaches its upper bound: $RRMSFE = 1$. If the uncertainty over the spot return is zero, $RRMSFE$ is ill-defined, but then also no hedging is needed since the amount of uncertainty cannot be reduced.

The hedging effectiveness measured by the $RRMSFE$ depends on the accuracy of the future spot price forecast, making the $RRMSFE$ a relative rather than an absolute measure. The best available forecast of the future spot price should be used to obtain a fair estimate of hedging effectiveness.

Caution is needed when comparing $RRMSFE$s across models. If the conditional mean forecasts of the future spot price, $\hat{s}_{t+h,w}$, differ across models, models with poorer spot price forecasts will by construction tend to yield higher hedging effectiveness. Therefore, $\hat{s}_{t+h,w}$ needs to be the same across models to allow for direct comparability of $RRMSFE$s. In the application I will obtain the spot price forecast $\hat{s}_{t+h,w}$ as an equally-weighted average of two elements: (1) a no-change forecast; (2) an equally-weighted forecast average from pairwise VECMs for the spot price paired with each hedging instrument/strategy under consideration. Averaging forecasts from different models should improve the forecast accuracy (Timmermann, 2006) and thus satisfy the need for the best available spot price forecast as mentioned above.

To the best of my knowledge, $RRMSFE$ has not been proposed as a measure of hedging effectiveness before. Even though the use of the $RRMSFE$ is justified by the particular objective function, we may want to consider it in the light of the classic measures. The golden standard in the hedging literature is Ederington's measure defined as relative reduction in variance (RRV) (Ederington, 1979). In the present context, Ederington's measure becomes

$$RRV = \frac{E_t((s_{t+h} - s_t)^2) - E_t\left(\left((s_{t+h} - \beta f_{t+h}) - (s_t - \beta f_t)\right)^2\right)}{E_t((s_{t+h} - s_t)^2)}. \tag{15}$$

Clearly, RRV is a special case of $RRMSFE$; when the expected spot return and the expected portfolio return are zero, $RRMSFE$ collapses to RRV. RRV may be a natural measure for applications in markets where neither asset returns nor portfolio returns are predictable, i.e. the best conditional mean prediction is the last observed value. However, RRV does not naturally accommodate cases of intrinsically predictable returns on individual assets and/or hedging portfolios as is the case with the salmon spot price. See Lien (2005) for an elaborate discussion of appropriateness (or lack thereof) of Ederington's measure in different contexts.

Whether measured by $RRMSFE$ or RRV, hedging effectiveness is to a large extent determined by the relevance of the hedging instrument. If the price of the hedging instrument co-moves closely with the price of the asset being hedged, hedging effectiveness will be high: $s_{t+h} - \beta f_{t+h}$ and $E_t(s_{t+h} - \beta f_{t+h})$ will be low and by the triangular inequality $E_t\left(\left((s_{t+h} - \beta f_{t+h}) - E_t(s_{t+h} - \beta f_{t+h})\right)^2\right)$ will be low, too, making the value of the $RRMSFE$ approach unity. Meanwhile, for a given asset and a fixed hedging instrument, hedging effectiveness will depend on two properties, the accuracy of the conditional mean forecast and the accuracy of the conditional variance forecast. First, accurate conditional mean forecasts of the spot price and the hedging instrument's price will yield accurate forecasts of a hedge portfolio for any given hedge ratio. The direction of this factor's effect on hedging efficiency is unclear; while accurate spot price forecasts alone reduce the $RRMSFE$, accurate forecasts of the portfolio value increase the $RRMSFE$. Second, the conditional variance forecasts determine the estimate of the optimal hedge ratio. Hence, the more accurate the conditional variance forecasts, the more accurate the estimates of the optimal hedge ratio and the higher the hedging effectiveness. (Note that assessing the forecast accuracy for the conditional variance is nontrivial as the realized variance is unobservable.) In conclusion, achieving high hedging effectiveness requires finding relevant hedging instruments and accurately predicting the conditional mean and variance of the spot price and the hedging instrument's price.

5 Data

The spot price of Norwegian farmed Atlantic salmon is obtained from NASDAQ[5]. Weekly survey data is available for 1995 and onwards, reflecting prices paid by exporters to salmon farmers (until 2013 week 13) and prices received by salmon exporters from foreign buyers (from 2013 week 14). I adjust for the difference between the earlier and the latter period by subtracting NOK 0.75/kg from the latter period's prices, following the practice of Fish Pool (Fish Pool, 2014). The spot price is available for seven different weight classes, from 1-2 kg though 6-7 kg and 7+ kg, and as a volume-weighted average of the seven prices. The average price together with its seasonal component and the seasonally adjusted average price are depicted in Figure 3.

Data for the FPI and the salmon futures prices is provided by Fish Pool[6]. Weekly FPI is available from 2006 onwards. The difference between the average spot price and the FPI was negligible in 2007-2015; the correlation between the two series was 0.997 in levels, 0.984 in the first differences and still as high as 0.989 in the second differences. Figure 4 depicts the FPI, original and seasonally adjusted, and its seasonal component.

Daily salmon futures prices are available from June 2006 onwards. I convert the prices from daily to weekly by taking the last price of each week. The front month salmon futures price is shown in Figure 5.

Daily share prices of Marine Harvest from September 1997 and daily Oslo Stock Exchange Benchmark Index (OSEBX) can be found at "Netfonds"[7]. I manually adjust Marine Harvest's share price for dividends and a reverse split, while OSEBX is adjusted by construction. I then convert the

---

[5] The data can be accessed at https://salmonprice.nasdaqomxtrader.com/public/home?1-1.ILinkListener-loginMenu-downloadIndexHistoryLink and https://salmonprice.nasdaqomxtrader.com/historicalNOSprices.xls

[6] The data can be accessed at http://fishpool.eu/price-information/spot-prices/history/ and http://fishpool.eu/fishpool/servlets/newt/public/excel?type=historic

[7] The data can be accessed at http://www.netfonds.no/quotes/paperdump.php?paper=MHG.OSE and http://www.netfonds.no/quotes/paperdump.php?paper=OSEBX.OSE

data from daily to weekly using Friday closing prices of each week; if Friday is not a trading day, the value from the preceding trading day is used. The data is plotted in Figure 6.

Weekly closing prices of live cattle, soybean meal and soybean oil futures are taken from “Quandl”[8]. The front month contracts’ prices are presented in Figures 7, 8 and 9.

Before modelling the joint behaviour of the salmon spot and the hedging instrument’s prices in detail, I carried out an initial assessment of relevance of the hedging instruments. I correlated the seasonally adjusted salmon spot prices with the seasonally adjusted hedging instruments’ prices, in levels and in first differences, both in the original currencies and in NOK. The correlations are presented in Table 1.

Figure 3 Salmon spot price: original (left), seasonally adjusted (middle), seasonal component (right), 2007 week 1 - 2015 week 17

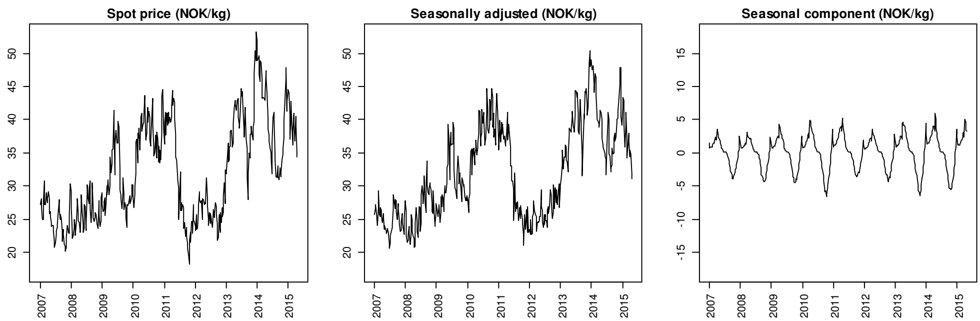

Figure 4 Fish Pool Index: original (left), seasonally adjusted (middle), seasonal component (right), 2007 week 1 - 2015 week 17

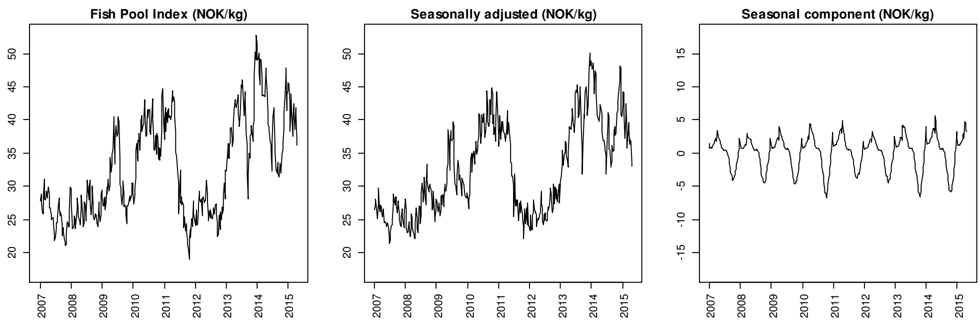

[8] The data can be accessed at https://www.quandl.com/collections/futures/cme-soybean-meal-futures, https://www.quandl.com/collections/futures/cme-soybean-oil-futures and https://www.quandl.com/collections/futures/cme-live-cattle-futures

Figure 5 Salmon front month futures price: original (left), seasonally adjusted (middle), seasonal component (right), 2007 week 1 - 2015 week 17

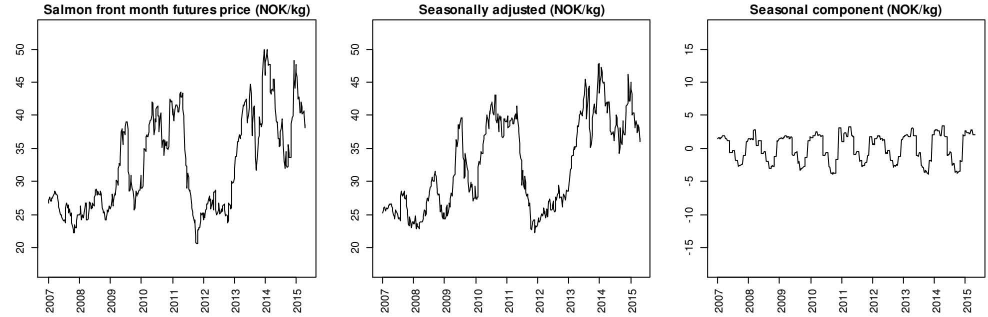

Figure 6 Marine Harvest share price (adjusted for dividend payments and a reverse split; left), Oslo Stock Exchange Benchmark Index (OSEBX; middle) and Marine Harvest share price adjusted for OSEBX (right), 2007 week 1 - 2015 week 17

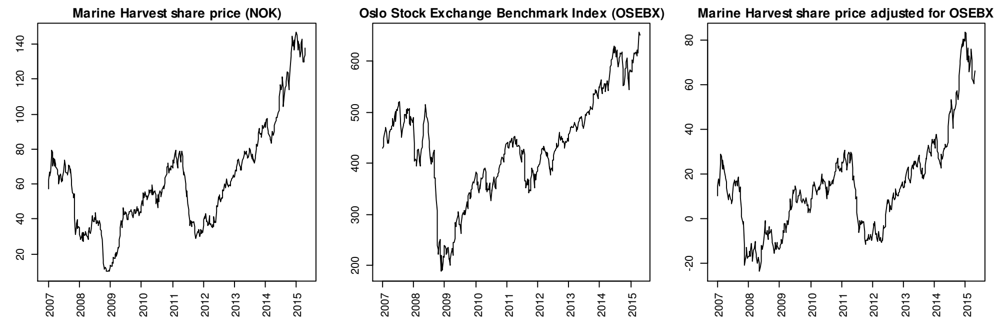

Figure 7 Live cattle front month futures price: original (left), seasonally adjusted (middle), seasonal component (right), 2007 week 1 - 2015 week 17

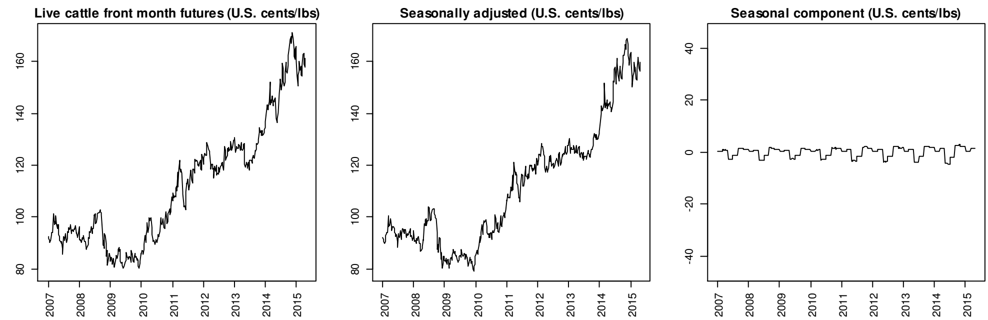

Figure 8 Soybean meal front month futures price: original (left), seasonally adjusted (middle), seasonal component (right), 2007 week 1 - 2015 week 17

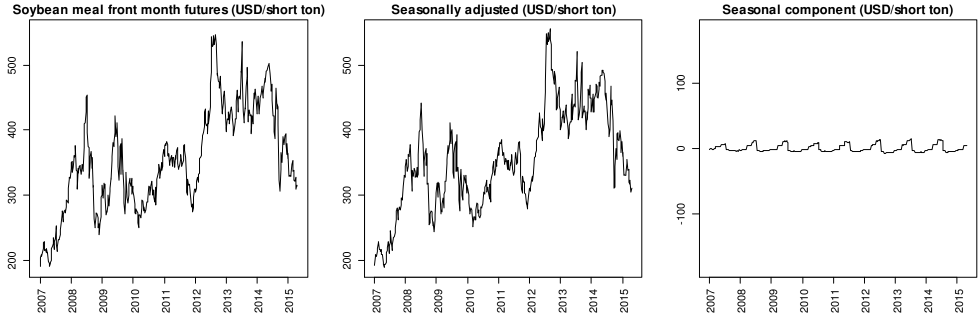

Figure 9 Soybean oil front month futures price: original (left), seasonally adjusted (middle), seasonal component (right), 2007 week 1 - 2015 week 17

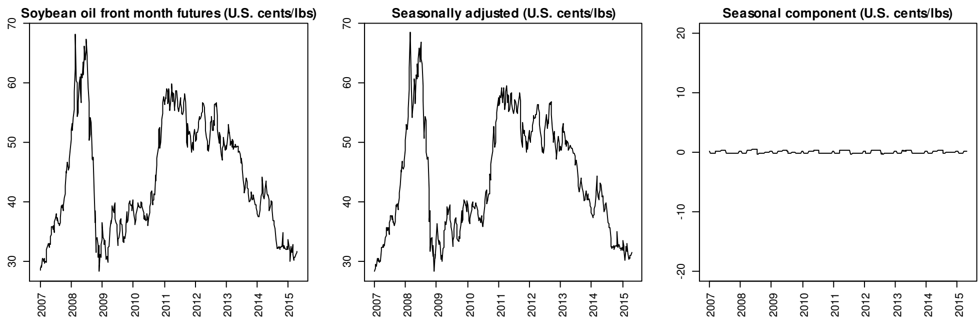

The higher the correlation, the more relevant the hedging instrument. Live cattle, soybean meal and soybean oil futures yielded level correlations of at most 0.45 and first difference correlations below 0.09 in absolute value in original currencies and marginally higher values when translated to NOK, which is relatively low. In the case of live cattle, apparently the substitution effect between salmon and beef is limited. As regards soybeans, it is only one of many production inputs; also, the short-term inelasticity of the salmon supply with respect to price has led to the spot price of salmon departing from the production costs considerably a number of times during the sample period; only in the long-term perspective may we expect a closer relation between the price of the production inputs and the salmon price. Meanwhile, Marine Harvest share price and adjusted share price had somewhat higher correlations with the salmon spot price, 0.61 in levels and up to 0.11 in the first differences. Considering all the hedging instruments but salmon futures yielded multiple correlation coefficients of up to 0.74 and 0.18 in absolute value. This still did not nearly match the salmon

Table 1 Correlations between seasonally adjusted average spot price and seasonally adjusted hedging instruments' prices, 2007 week 1 - 2015 week 17

	Levels, original currency	First differences, original currency	Levels, NOK	First differences, NOK
Salmon futures, front month	**0.96**	**0.39**		
Salmon futures, next month	**0.93**	**0.35**		
Live cattle futures, front month	0.45	0.03	0.48	0.11
Live cattle futures, next month	0.43	0.08	0.47	0.15
Soybean meal futures, front month	0.37	-0.04	0.50	0.00
Soybean meal futures, next month	0.32	-0.06	0.46	-0.01
Soybean oil futures, front month	-0.21	-0.07	-0.09	-0.01
Soybean oil futures, next month	-0.22	-0.08	-0.10	-0.01
Marine Harvest share	0.61	-0.01		
Marine Harvest share, adjusted for OSEBX	0.61	0.11		
All but salmon futures*	0.74	0.16	0.71	0.18

* Multiple correlation, obtained as the square root of the coefficient of determination (R^2) from a multiple regression of spot price on hedging instruments' prices

futures contracts with correlations above 0.93 in levels and 0.35 in the first differences. Seeing the clear superiority of the salmon futures and limited potential for effective hedging with the remaining instruments due to their low correlations with the spot price of salmon, in the remainder I will only pursue hedging with salmon futures.

The higher the correlation, the more relevant the hedging instrument. Live cattle, soybean meal and soybean oil futures yielded level correlations of at most 0.45 and first difference correlations below 0.09 in absolute value in original currencies and marginally higher values when translated to NOK, which is relatively low. In the case of live cattle, apparently the substitution effect between salmon and beef is limited. As regards soybeans, it is only one of many production inputs; also, the short-term inelasticity of the salmon supply with respect to price has led to the spot price of salmon departing from the production costs considerably a number of times during the sample period; only in the long-term perspective may we expect a closer relation between the price of the production inputs and the salmon price. Meanwhile, Marine Harvest share price and adjusted share price had somewhat higher correlations with the salmon spot price, 0.61 in levels and up to 0.11 in the first differences. Considering all the hedging instruments but salmon futures yielded multiple correlation coefficients of up to 0.74 and 0.18 in absolute value. This still did not nearly match the salmon futures contracts with correlations above 0.93 in levels and 0.35 in the first differences. Seeing the clear superiority of the salmon futures and limited potential for effective hedging with the remaining instruments due to their low correlations with the spot price of salmon, in the remainder I will only pursue hedging with salmon futures.

6 Results and discussion

Out-of-sample hedging results are presented in this chapter. The assets being hedged are the seven weight classes of salmon and the average salmon sold in the spot market. The only hedging instrument considered is the salmon futures employed using three different strategies, F0, F1 and M, as described in the chapter 3. There are four models used for predicting the conditional mean and variance: (1) no-change combined with constant variance; (2) no-change combined with DCC-GARCH; (3) VECM combined with constant variance; and (4) VECM combined with DCC-GARCH. Hedging effectiveness is measured for three horizons: four, eight and 13 weeks ahead.

6.1 Hedging effectiveness

Hedging effectiveness as measured by $\widehat{RRMSFE}$ is given in Tables 2.A through 2.D and the corresponding Figure 10. The values range from 0.07 to 0.60 indicating a 7% to 60% reduction in the mean squared forecast error when using a hedge portfolio in place of an unhedged spot position.

The most effective hedging strategy is M, regardless of the hedging horizon. F0 is the worst for the 4-week hedge while F1 is the worst for the 8-week and the 13-week hedges. Hence, keeping the futures contract through maturity is not only the cheapest but also the most effective strategy.

The VEC+DCC-GARCH model yields the highest hedging effectiveness and is followed by the no-change+DCC-GARCH model. VEC conditional mean model with constant conditional variance performs the worst. The differences in the $\widehat{RRMSFE}$ between models with different conditional mean specifications but the same conditional variance specification are small, on average about 0%p-2%p (%p denotes percentage points). Hence, choosing a VECM in place of a no-change model might not increase hedging effectiveness significantly. However, the differences between models with the same conditional mean specification but different conditional variance specifications are considerable, on average about 3%p-7%p. Thus choosing a DCC-GARCH conditional variance specification instead of constant conditional variance brings a substantial increase in hedging effectiveness.

Hedging effectiveness increases with the hedging horizon, as should be expected when the spot and the futures prices are cointegrated. The cumulative effect of the error correction mechanism bringing the cointegrated series to a fixed equilibrium soars with time and marginalizes the effect of the short term fluctuations due to random price shocks. In other words, if two time series share a

Table 2.A Relative reduction in mean squared forecast error ($\widehat{RRMSFE}$) of hedge portfolios as compared to unhedged position for no-change conditional mean and constant conditional variance model, 2007 week 1 - 2015 week 17

Spot price / Hedging instrument	1-2 kg	2-3 kg	3-4 kg	4-5 kg	5-6 kg	6-7 kg	7+ kg	Avg.
4 weeks ahead								
F0	0.07	0.13	0.10	0.12	0.13	0.13	0.12	0.12
F1	0.11	0.18	0.18	0.19	0.18	0.18	0.18	0.18
M	0.14	0.21	0.23	0.24	0.23	0.22	0.21	0.23
8 weeks ahead								
F0	0.28	0.38	0.36	0.36	0.34	0.35	0.33	0.36
F1	0.32	0.40	0.38	0.36	0.31	0.33	0.29	0.36
M	0.29	0.39	0.39	0.39	0.37	0.37	0.34	0.39
13 weeks ahead								
F0	0.40	0.53	0.48	0.47	0.43	0.45	0.43	0.47
F1	0.38	0.49	0.43	0.41	0.37	0.37	0.36	0.41
M	0.41	0.54	0.54	0.54	0.50	0.51	0.49	0.54

Table 2.B Relative reduction in mean squared forecast error ($\widehat{RRMSFE}$) of hedge portfolios as compared to unhedged position for no-change conditional mean and DCC-GARCH conditional variance model, 2007 week 1 - 2015 week 17

Spot price / Hedging instrument	1-2 kg	2-3 kg	3-4 kg	4-5 kg	5-6 kg	6-7 kg	7+ kg	Avg.
4 weeks ahead								
F0	0.11	0.18	0.12	0.16	0.13	0.14	0.14	0.16
F1	0.23	0.26	0.25	0.26	0.21	0.22	0.21	0.26
M	0.21	0.25	0.24	0.26	0.24	0.23	0.23	0.24
8 weeks ahead								
F0	0.35	0.42	0.41	0.42	0.39	0.39	0.39	0.42
F1	0.34	0.40	0.38	0.39	0.35	0.35	0.34	0.39
M	0.37	0.42	0.42	0.45	0.41	0.41	0.40	0.44
13 weeks ahead								
F0	0.44	0.54	0.53	0.53	0.49	0.49	0.48	0.52
F1	0.39	0.48	0.46	0.44	0.41	0.41	0.40	0.44
M	0.47	0.55	0.58	0.58	0.54	0.53	0.53	0.57

Table 2.C Relative reduction in mean squared forecast error ($\widehat{\boldsymbol{RRMSFE}}$) of hedge portfolios as compared to unhedged position for VEC conditional mean and constant conditional variance model, 2007 week 1 - 2015 week 17

Spot price / Hedging instrument	1-2 kg	2-3 kg	3-4 kg	4-5 kg	5-6 kg	6-7 kg	7+ kg	Avg.
4 weeks ahead								
F0	0.08	0.12	0.11	0.13	0.13	0.12	0.11	0.12
F1	0.11	0.16	0.18	0.19	0.17	0.16	0.15	0.18
M	0.16	0.21	0.24	0.25	0.24	0.21	0.20	0.23
8 weeks ahead								
F0	0.27	0.37	0.36	0.35	0.33	0.33	0.31	0.35
F1	0.31	0.35	0.36	0.34	0.30	0.29	0.26	0.34
M	0.28	0.36	0.37	0.38	0.35	0.34	0.33	0.37
13 weeks ahead								
F0	0.38	0.50	0.48	0.48	0.44	0.44	0.43	0.48
F1	0.35	0.44	0.41	0.40	0.35	0.36	0.35	0.40
M	0.39	0.51	0.52	0.53	0.49	0.48	0.48	0.52

Table 2.D Relative reduction in mean squared forecast error ($\widehat{\boldsymbol{RRMSFE}}$) of hedge portfolios as compared to unhedged position for VEC conditional mean and DCC-GARCH conditional variance model, 2007 week 1 - 2015 week 17

Spot price / Hedging instrument	1-2 kg	2-3 kg	3-4 kg	4-5 kg	5-6 kg	6-7 kg	7+ kg	Avg.
4 weeks ahead								
F0	0.10	0.18	0.12	0.17	0.15	0.16	0.15	0.17
F1	0.24	0.26	0.28	0.28	0.24	0.23	0.22	0.29
M	0.22	0.26	0.26	0.29	0.27	0.24	0.24	0.27
8 weeks ahead								
F0	0.36	0.41	0.42	0.44	0.41	0.41	0.40	0.43
F1	0.35	0.37	0.39	0.40	0.37	0.37	0.34	0.40
M	0.39	0.41	0.43	0.46	0.44	0.43	0.42	0.45
13 weeks ahead								
F0	0.45	0.53	0.54	0.56	0.53	0.52	0.51	0.54
F1	0.39	0.45	0.45	0.45	0.43	0.42	0.41	0.44
M	0.49	0.54	0.58	0.60	0.58	0.56	0.56	0.58

common stochastic trend, in the long run their paths will be roughly the same despite any short-term deviations from the equilibrium due to stationary zero-mean random shocks.

There is no clear pattern to the variation in the hedging effectiveness across the different weight classes. Only the smallest 1-2 kg weight class stands out with poor hedging performance. This matches the finding in Bloznelis (2016a) that the price dynamics of the 1-2 kg fish is the most distinct among the weight classes.

6.2 Forecasting performance

As mentioned in the chapter 3, hedging effectiveness is determined by the forecast accuracy of the conditional mean and variance models, although only the former can be easily assessed. I will now examine the forecasting performance of the conditional mean models looking for explanations for the variation in the hedging effectiveness across the weight classes, hedging strategies, models and hedging horizons.

The forecast accuracy of the different methods is given in Tables 3.A through 3.D in terms of root mean squared forecast error ($RMSFE$). Mean squared forecast error could be used instead so as to match the hedger's objective function; however, $RMSFE$ provides a more convenient interpretation as its scale corresponds to the scale of the original series. $RMSFE$ lies in the range between NOK 2.8-6.1/kg. Naturally, the shorter the forecast horizon, the more accurate the forecasts.

From the weight class perspective, $RMSFE$ increases with fish size; the 5-7+ kg weight classes are more difficult to predict than the smaller fish. However, this effect is mostly due to scale as smaller fish are cheaper than large ones. Once $RMSFE$ is scaled by the average price for each weight class, the pattern largely disappears. That is in line with the lack of regularity in the variation of hedging effectiveness across the weight classes (note that the measure of hedging effectiveness, $RRMSFE$, is not sensitive to scale).

The portfolios in which the futures contract is held until maturity (the M portfolios) are forecasted the most accurately among the three alternative strategies regardless of the length of the forecast horizon. F1 works better than F0 in the 4-week horizon, but that is reversed for the 8-week and 13-week horizons. This is exactly the same picture as in the case of hedging effectiveness. Thus the ability to forecast the returns on the M portfolios better than the returns on F0 and F1 portfolios is among the causes of the superior hedging effectiveness of the M strategy.

The best forecasting model is the VECM with DCC-GARCH errors, followed by the VECM with constant conditional error variance. The no-change forecast with constant error variance and the no-change forecast with DCC-GARCH errors provide less accurate forecasts. However, the differences in forecast accuracy across the models are small, up to NOK 0.5/kg. Let us compare the variation in

Figure 10 Relative reduction in mean squared forecast error ($\widehat{RRMSFE}$) of hedge portfolios as compared to unhedged position for the different models, hedging horizons, weight classes and strategies, 2007 week 1 - 2015 week 17

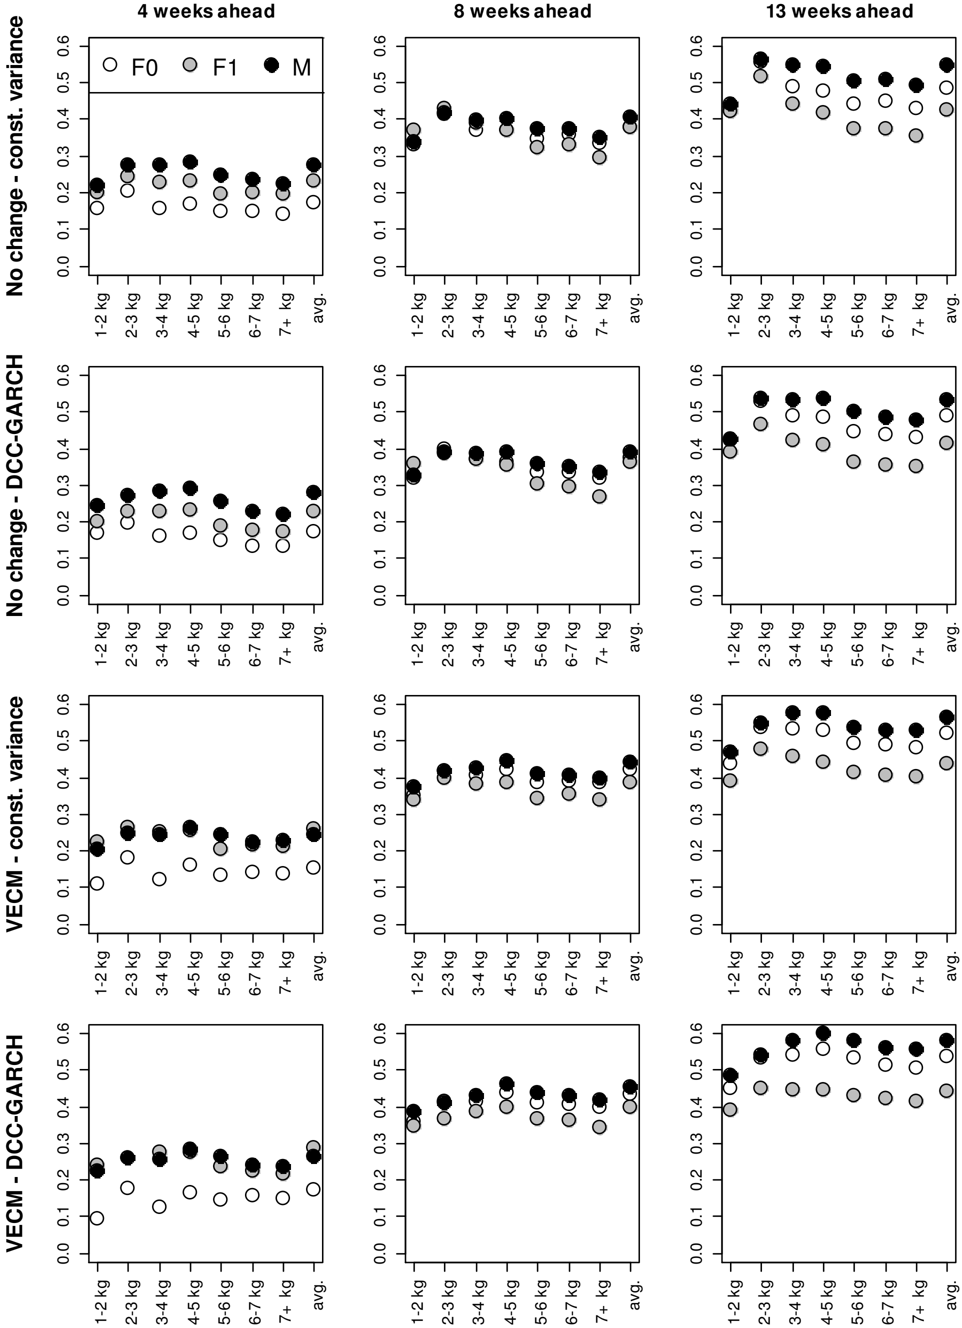

Table 3.A Root mean squared forecast error ($RMSFE$) of hedge portfolios for no-change conditional mean and constant conditional variance model, 2007 week 1 - 2015 week 17

Spot price / Hedging instrument	1-2 kg	2-3 kg	3-4 kg	4-5 kg	5-6 kg	6-7 kg	7+ kg	Avg.
<u>4 weeks ahead</u>								
Unhedged spot	3.17	3.55	3.81	3.97	4.44	4.45	4.49	3.89
F0	3.06	3.31	3.61	3.72	4.15	4.15	4.21	3.66
F1	2.98	3.22	3.46	3.58	4.03	4.03	4.08	3.52
M	2.94	3.16	3.35	3.45	3.90	3.94	4.01	3.42
<u>8 weeks ahead</u>								
Unhedged spot	3.81	4.52	4.89	5.21	5.86	6.04	6.07	5.11
F0	3.24	3.55	3.92	4.17	4.77	4.87	4.98	4.09
F1	3.14	3.50	3.86	4.17	4.86	4.97	5.13	4.09
M	3.22	3.55	3.83	4.07	4.67	4.81	4.93	4.00
<u>13 weeks ahead</u>								
Unhedged spot	4.67	5.73	6.13	6.52	7.04	7.53	7.60	6.39
F0	3.63	3.92	4.43	4.75	5.30	5.60	5.74	4.64
F1	3.67	4.10	4.64	5.01	5.61	5.97	6.10	4.90
M	3.61	3.88	4.16	4.42	4.98	5.29	5.42	4.35

Table 3.B Root mean squared forecast error ($RMSFE$) of hedge portfolios for no-change conditional mean and DCC-GARCH conditional variance model, 2007 week 1 - 2015 week 17

Spot price / Hedging instrument	1-2 kg	2-3 kg	3-4 kg	4-5 kg	5-6 kg	6-7 kg	7+ kg	Avg.
<u>4 weeks ahead</u>								
Unhedged spot	3.17	3.55	3.81	3.97	4.44	4.45	4.49	3.89
F0	3.04	3.33	3.61	3.72	4.14	4.19	4.24	3.66
F1	2.99	3.26	3.46	3.58	4.05	4.08	4.14	3.53
M	2.90	3.16	3.34	3.43	3.87	3.96	4.02	3.41
<u>8 weeks ahead</u>								
Unhedged spot	3.81	4.52	4.89	5.21	5.86	6.04	6.07	5.11
F0	3.26	3.60	3.92	4.20	4.82	4.96	5.04	4.11
F1	3.17	3.64	3.92	4.23	4.92	5.10	5.22	4.15
M	3.25	3.63	3.87	4.12	4.73	4.90	4.99	4.05
<u>13 weeks ahead</u>								
Unhedged spot	4.67	5.73	6.13	6.52	7.04	7.53	7.60	6.39
F0	3.67	4.03	4.44	4.71	5.27	5.66	5.74	4.62
F1	3.77	4.30	4.72	5.05	5.66	6.05	6.12	4.95
M	3.66	4.01	4.25	4.46	5.01	5.41	5.49	4.41

Table 3.C Root mean squared forecast error ($RMSFE$) of hedge portfolios for VEC conditional mean and constant conditional variance model, 2007 week 1 - 2015 week 17

Spot price / Hedging instrument	1-2 kg	2-3 kg	3-4 kg	4-5 kg	5-6 kg	6-7 kg	7+ kg	Avg.
4 weeks ahead								
Unhedged spot	3.17	3.55	3.81	3.97	4.44	4.45	4.49	3.89
F0	2.99	3.22	3.58	3.64	4.13	4.12	4.17	3.58
F1	2.79	3.05	3.30	3.43	3.96	3.94	3.99	3.35
M	2.82	3.08	3.32	3.41	3.86	3.92	3.95	3.38
8 weeks ahead								
Unhedged spot	3.81	4.52	4.89	5.21	5.86	6.04	6.07	5.11
F0	3.08	3.45	3.77	3.97	4.59	4.72	4.76	3.89
F1	3.10	3.51	3.84	4.09	4.75	4.86	4.94	4.00
M	3.02	3.45	3.71	3.88	4.51	4.66	4.70	3.82
13 weeks ahead								
Unhedged spot	4.67	5.73	6.13	6.52	7.04	7.53	7.60	6.39
F0	3.50	3.89	4.18	4.47	5.01	5.39	5.47	4.43
F1	3.66	4.14	4.52	4.86	5.40	5.80	5.88	4.79
M	3.40	3.85	4.00	4.25	4.78	5.16	5.23	4.21

Table 3.D Root mean squared forecast error ($RMSFE$) of hedge portfolios for VEC conditional mean and DCC-GARCH conditional variance model, 2007 week 1 - 2015 week 17

Spot price / Hedging instrument	1-2 kg	2-3 kg	3-4 kg	4-5 kg	5-6 kg	6-7 kg	7+ kg	Avg.
4 weeks ahead								
Unhedged spot	3.17	3.55	3.81	3.97	4.44	4.45	4.49	3.89
F0	3.01	3.22	3.57	3.63	4.11	4.09	4.15	3.54
F1	2.76	3.05	3.25	3.38	3.88	3.92	3.98	3.28
M	2.79	3.06	3.29	3.36	3.81	3.88	3.93	3.34
8 weeks ahead								
Unhedged spot	3.81	4.52	4.89	5.21	5.86	6.04	6.07	5.11
F0	3.05	3.46	3.74	3.91	4.50	4.65	4.71	3.85
F1	3.08	3.59	3.83	4.04	4.66	4.82	4.93	3.96
M	2.99	3.48	3.69	3.82	4.40	4.56	4.63	3.78
13 weeks ahead								
Unhedged spot	4.67	5.73	6.13	6.52	7.04	7.53	7.60	6.39
F0	3.46	3.92	4.15	4.34	4.82	5.25	5.34	4.35
F1	3.65	4.25	4.56	4.84	5.32	5.73	5.82	4.77
M	3.35	3.88	3.96	4.11	4.57	4.99	5.06	4.13

hedging effectiveness and forecast accuracy across the different models for a fixed strategy. While the forecast accuracy of the conditional mean does not vary much with the choice of the conditional variance model, the hedging effectiveness does. Therefore, it is the forecast accuracy of the conditional variance model that determines the hedging effectiveness to a considerable degree. Similarly, let us examine the variation in hedging effectiveness and forecast accuracy across the different strategies for a fixed model. The variation in hedging effectiveness across the hedging strategies largely mimics the variation in forecast accuracy. If the variation in hedging effectiveness may be explained by forecast accuracy, then we do not have evidence that hedging effectiveness is determined by the relevance of the hedging strategy itself (apart from the fact that outcomes of different hedging strategies have different degrees of predictability). These findings should be viewed as indicative only; meanwhile, disentangling the different factors' effects on hedging effectiveness quantitatively may be prohibitively complicated and is not pursued in this study.

6.3 Optimal hedge ratios

When hedging the uncertainty in the salmon spot price with a salmon futures contract, the hedge ratio gives the exposure to the futures price as a fraction (or a multiple) of the exposure to the spot price. In other words, it is a ratio of the financial exposure to the physical exposure. A hedge ratio below one indicates the financial exposure is less than the physical exposure, while a hedge ratio above one shows the opposite. The optimal hedge ratios for the average spot price of salmon produced by the VEC+DCC-GARCH model are depicted in Figure 11. Each point in a given panel of the figure corresponds to a hedge ratio obtained from a different rolling window. E.g. the first point in the top left graph is dated 2011 week 4 and shows the 4-week-ahead hedge ratio obtained from a rolling window spanning from 2007 week 1 through 2010 week 52; the second point is dated 2011 week 5 and shows the 4-week-ahead hedge ratio obtained from a rolling window spanning from 2007 week 2 through 2011 week 1; etc.

The variation within one curve reflects the variation across the rolling windows. It is quite high for the 4-week-ahead hedge but decreases with the hedging horizon, as was to be expected. The h-week hedge ratio is the ratio of an h-element sum of predicted covariances over an h-element sum of predicted variances as per equation (12). For the hedging horizon $h = 4$, one element out of four changes in each of the sums when moving from one rolling window to the next one. Meanwhile, for a 13-week hedge only one element out of thirteen changes in each of the sums. Thus naturally there is more variation in the 4-week than in the 13-week hedge ratio.

From a longer perspective, we see that the hedge ratios for F0 and F1 (and to a small extent also for M) were higher in the beginning of the sample but decreased over the first two years, representing the rolling windows from 2007-2010 through 2009-2012. There has been no clear trend

Figure 11 Optimal hedge ratios from VEC+DCC-GARCH model for hedging the average spot price of salmon, 2011 week 1 - 2015 week 17. Horizontal lines mark 0, 1 and the mean of the optimal hedge ratio over the period

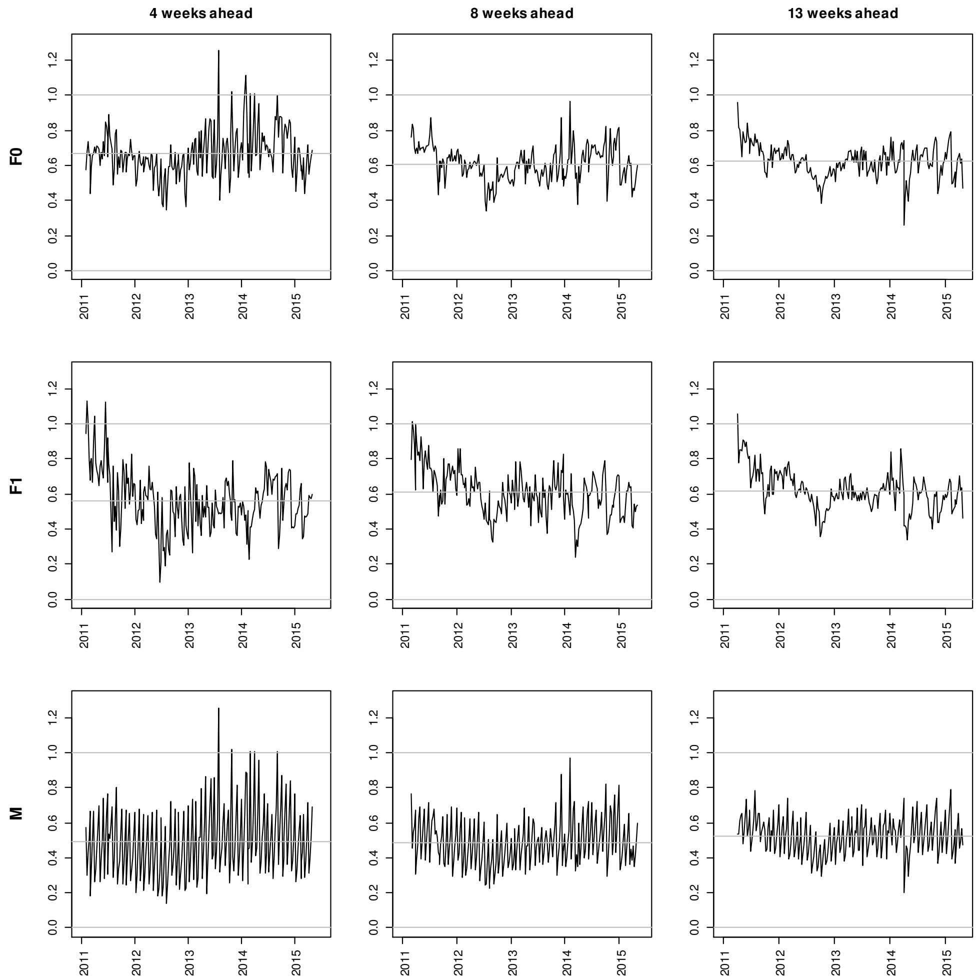

afterwards. Examining the predicted covariances and variances as in equation (12) more closely does not allow drawing broad conclusions as to what the driving force was, since there have been considerable variations across the weight classes and the hedging strategies.

There are also visible differences across the hedging strategies. The average hedge ratios are higher for F0 and F1 (around 0.6) than for M (around 0.5). There are substantial high-frequency oscillations in the case of M, which can be explained by the nature of keeping the futures contract until maturity. Consider hedging a spot position that is in week 1 (the first week) of a month.

Following the M strategy, the front month futures contract will be kept for another three or four weeks until expiration. The spot and the futures prices will both move only in week 1; after that, only the futures price will move as the spot position will have already been liquidated. As per equation (12), the sum of the predicted covariances in the numerator will contain only one non-zero element. Meanwhile, the sum of the predicted variances in the denominator will have four or five non-zero elements. Thus the value of the fraction, which is the hedge ratio, will be relatively low in absolute value. Now consider hedging a spot position that is in the last week of a month. The sum of the predicted covariances in the numerator will now contain four or five non-zero elements, while the denominator will stay the same as before. Since the predicted covariances are likely to be positive for each hedging horizon, their four- or five-element sum will be greater than just its first element. Thus, when hedging the last week of a month, the fraction representing the hedge ratio will be relatively large. For the weeks in between (week 2 and week 3 in all months, and week 4 in 5-week months) the effect will be in between these two cases. Thus the high-frequency oscillations with periods of four or five weeks are as expected.

The optimal hedge ratios for the specific weight classes are not presented due to space limitations, but their patterns are quite similar to the ones for the average spot price observed in Figure 11. However, the smallest fish have a lower average hedge ratio than the remaining weight classes. Also, the optimal hedge ratios for models other than the VECM+DCC-GARCH are skipped for brevity.

6.4 How attractive is hedging in practical terms?

How attractive does hedging with salmon futures appear in practical terms? The 4-week hedging effectiveness of the most successful M strategy yields $R\widehat{RMSFE}$ of 23-27% for the average spot price. For ease of interpretation, let us view this in terms of root mean squared forecast error ($RMSFE$); the reduction is from an $RMSFE$ of NOK 3.9/kg to NOK 3.3/kg, which does not seem substantial. Adding the poor liquidity of the futures market and the transaction costs (although quite small for the M strategy), the 4-week hedge is unlikely to attract considerable interest. Meanwhile, $R\widehat{RMSFE}$ of 52-58% for the 13-week hedge, or a reduction from NOK 6.4/kg to NOK 4.1/kg in $RMSFE$, seems more attractive. Also, recall that hedging effectiveness would likely increase if monthly rather than weekly spot positions were hedged, and nearly perfect hedges could be expected. This makes hedging with salmon futures even more appealing. Why then are there so few trades at Fish Pool? The possible reasons could be, (1) the salmon farmers' insensitivity to risk (Bergfjord, 2009), (2) lack of speculative interest and thus lack of counterparties to the potential hedgers, (3) increasing vertical integration in the salmon industry providing implicit hedging within a vertically integrated company, and (4) lack of independent studies assessing the hedging effectiveness and providing reliable and timely information; this gap is partly filled with the current study.

6.5 Conclusion

The uncertainty in the future spot price of salmon and ways of reducing it has been investigated in the recent papers by Bloznelis (Bloznelis, 2016a, Bloznelis, 2016b) and in the current manuscript, but the problem appears to be a hard nut to crack. Theoretically, hedging the price uncertainty with the salmon futures could be regarded as perhaps the most efficient solution. However, the lack of liquidity in the futures market for salmon limits the practical applicability of this option. Unless the liquidity picks up in the future, the salmon industry will have to rely on the ongoing vertical integration, which provides implicit hedging within a company, and perhaps on increased use of salmon forward contracts. This offers two venues for future research, one on the nature and effectiveness of the implicit within-company hedging and another on the analysis of the forward market for salmon.

7 Acknowledgements

I thank Ole Gjølberg, Ronald Huisman, Marius Bausys, Arnar Mar Buason and members of the Norwegian Center for Commodity Market Analysis for valuable discussions and feedback. All remaining errors are mine.

References

ASCHE, F. & BJORNDAL, T. 2011. *The economics of salmon aquaculture,* Chichester, John Wiley & Sons.

ASCHE, F. & GUTTORMSEN, A. G. 2001. Patterns in the relative price for different sizes of farmed fish. *Marine Resource Economics,* 16**,** 235-247.

BERGFJORD, O. J. 2007. Is there a future for salmon futures? An analysis of the prospects of a potential futures market for salmon. *Aquaculture Economics & Management,* 11**,** 113-132.

BERGFJORD, O. J. 2009. Risk perception and risk management in Norwegian aquaculture. *Journal of Risk Research,* 12**,** 91-104.

BLOZNELIS, D. 2016a. Salmon price volatility: a weight-class-specific multivariate approach. *forthcoming in Aquaculture Economics & Management*.

BLOZNELIS, D. 2016b. Short-term salmon price forecasting. *Preprint.*

BOLLERSLEV, T. 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics,* 31**,** 307-327.

CARREIRA, R. M. 2013. *The causal relationship between salmon prices and share price (2013). Price analysis on the Oslo Stock Exchange.* Masters, The Arctic University of Norway.

CHEN, S.-S., LEE, C.-F. & SHRESTHA, K. 2003. Futures hedge ratios: a review. *The Quarterly Review of Economics and Finance,* 43**,** 433-465.

EDERINGTON, L. H. 1979. The hedging performance of the new futures markets. *The Journal of Finance,* 34**,** 157-170.

ENGLE, R. 2002. Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics,* 20**,** 339-350.

FISH POOL. 2014. *NOS and NASDAQ price* [Online]. Fish Pool. Available: http://fishpool.eu/wp-content/uploads/2014/08/FAQ-Nasdaq.pdf [Accessed 2015-03-30].

FISH POOL 2015. Arsrapport 2014. Oslo/Bergen/Florø.

GARCIA, P. & LEUTHOLD, R. M. 2004. A selected review of agricultural commodity futures and options markets. *European Review of Agricultural Economics,* 31**,** 235-272.

GUTTORMSEN, A. G. 2015-01-22 2015. Type to BLOZNELIS, D.

HANSEN, P. R. & LUNDE, A. 2005. A forecast comparison of volatility models: does anything beat a GARCH (1, 1)? *Journal of Applied Econometrics,* 20**,** 873-889.

HYNDMAN, R. J. 2014. Forecasting weekly data. *Hyndsight* [Online]. Available from: http://robjhyndman.com/hyndsight/forecasting-weekly-data/ [Accessed 2014-03-05].

HULL, J. C. 2012. *Options, futures, and other derivatives,* London, Pearson Education Limited.

INVESTOPEDIA. *Hedge* [Online]. Available: http://www.investopedia.com/terms/h/hedge.asp [Accessed 2015-09-16].

JENSEN, B.-A. 2013. *No immediate relief in sight for high salmon prices* [Online]. Available: http://www.intrafish.com/news/article1373426.ece [Accessed 2014-02-28].

KONISHI, S. & KITAGAWA, G. 2008. *Information criteria and statistical modeling,* New York, Springer Science & Business Media.

LARSEN, T. A. & ASCHE, F. 2011. Contracts in the salmon aquaculture industry: An analysis of Norwegian salmon exports. *Marine Resource Economics,* 26**,** 141-150.

LENCE, S. H. 1995. The economic value of minimum-variance hedges. *American Journal of Agricultural Economics,* 77**,** 353-364.

LENCE, S. H. 1996. Relaxing the assumptions of minimum-variance hedging. *Journal of Agricultural and Resource Economics,* 21**,** 39-55.

LIEN, D. 2005. The use and abuse of the hedging effectiveness measure. *International Review of Financial Analysis,* 14**,** 277-282.

LINTNER, J. 1965. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics,* 47**,** 13-37.

MARINE HARVEST 2015. Salmon farming industry handbook 2015. Bergen: Marine Harvest.

OGLEND, A. & ASCHE, F. 2015. Salmon and Feed Raw Material Markets. *Preprint*.

SHARPE, W. F. 1964. Capital asset prices: a theory of market equilibrium under conditions of risk. *The Journal of Finance,* 19**,** 425-442.

TIMMERMANN, A. 2006. Forecast combinations. *In:* ELLIOTT, G., GRANGER, C. W. J. & TIMMERMANN, A. (eds.) *Handbook of Economic Forecasting.* 1 ed.

Daumantas Bloznelis

Daumantas Bloznelis was born in Vilnius, Lithuania, in 1986. He holds BSc (2009) and MSc (2011) degrees in Statistics/Econometrics from Vilnius University, Lithuania.

The thesis deals with the uncertainty over the future spot price of salmon. It consists of an introduction and three independent papers. The introduction provides an overview of the salmon farming industry and the salmon market, lists the research problems and objectives, presents methods, software implementation and data sources, and summarizes the findings, contributions and limitations.

Paper I describes salmon price volatility using ARMA and DCC-GARCH models. Two periods of different volatility regimes are identified, before and after 2006 when the salmon market was undergoing fundamental changes. Both volatility and conditional correlations increase in the second period, and return dynamics become more homogenous across different weight classes of salmon. This development is conducive to the functioning of the salmon futures and options exchange.

Paper II deals with short-term salmon price forecasting. It employs classical time series models with and without regularization (shrinkage), artificial neural networks and the k-nearest neighbours method. Salmon price is found to be essentially unpredictable beyond a seasonal component, which indicates weak form efficiency in the salmon spot market.

Paper III concerns hedging the uncertainty in the future spot price of salmon using salmon futures, soybean and live cattle futures and share prices of salmon farming companies. A new objective function for hedging and a new measure of hedging effectiveness are proposed. The only instrument to offer considerable hedging effectiveness is the salmon futures contracts. Interestingly, holding the contract through maturity is not only cheaper but also more efficient than liquidating the futures and the spot positions simultaneously.

In conclusion, hedging emerges as the preferred option for uncertainty management, provided there is enough liquidity in the salmon futures market. Among the alternatives not explored in the thesis, forward trading could be considered. Also, the uncertainty problem may subside over time due to the ongoing vertical integration in the salmon industry.

Professors Atle Guttormsen and Ole Gjølberg were Daumantas' supervisors

E-mail: dbloznelis@gmail.com

School of Economics and Business
Norwegian University of Life Sciences (NMBU)
P.O Box 5003
N-1432 Ås, Norway

Telephone: +47 67231100
e-mail*: hh@nmbu.no*
http:/www.nmbu.no/hh

ISSN: 1894-6402
ISBN: 978-82-575-1377-1