Dynamic Factor Modelling of Major Norwegian Export Commodities

Decomposing and forecasting the price movements of Atlantic salmon and Brent Crude oil

Sigve Kristiansen Borgmo
Magnus Aabø Mellem
Mads Sætre Michalsen

Industrial Economics and Technology Management
Submission date: June 2018
Supervisor: Petter Eilif de Lange, IØT

Norwegian University of Science and Technology
Department of Industrial Economics and Technology Management

Problem Description

The objective of this thesis is to apply *dynamic factor modelling* to the study of two commodities highly material to the Norwegian economy. We contribute to the literature in the following ways:

Co-movement and composition analysis - Considering a large data set of globally traded commodities, we apply a dynamic factor model (DFM) approach in order to estimate the level of co-movement in global commodity markets, and decompose individual commodity price returns into global, sectoral, and commodity-specific components. Specifically, we estimate the share of individual commodities' price movements attributable to each of the following sources: 1) broader, common trends in global commodity markets and the global business cycle, 2) common trends within sub-groups of related commodities, and 3) shocks unique and isolated to the individual commodity markets. We focus our analysis on two major Norwegian export commodities, Atlantic salmon and Brent Crude oil.

Spot price forecasting - We assess the DFM framework's applicability to the task of forecasting monthly spot prices for Atlantic salmon and Brent Crude oil. There exists no previous studies dedicated to forecasting these commodities using a DFM approach, so with our study, new ground is broken.

Preface

This thesis, written within the field of Financial Engineering, concludes our Master of Science degrees in Industrial Economics and Technology Management at the Norwegian University of Science and Technology (NTNU). It is original and independent work performed by Sigve K. Borgmo, Magnus A. Mellem and Mads S. Michalsen, written in the spring semester of 2018.

We would like to thank our supervisor, Associate Professor Petter E. de Lange at the Department of International Business at NTNU, for valuable advice and guidance. His interest in our work has been truly helpful and motivating during the process of writing this thesis. A word of appreciation is also offered to our friends and families, several of whom have contributed with both insightful feedback with regards to this thesis, as well as invaluable support throughout our time at NTNU.

Abstract

This thesis studies the price movements of two major Norwegian export commodities, Atlantic salmon and Brent Crude oil, using a dynamic factor modelling approach. We pursue two different avenues of research, namely *co-movement and composition analysis* of commodity price returns and *forecasting* of future spot prices.

In the first study, we investigate dynamic co-movement in global commodity markets via a restricted dynamic factor model (DFM) and decompose commodity returns into global, sectoral, and idiosyncratic components. This decomposition allows us to estimate the degree to which global, sectoral and commodity-specific factors explain the fluctuations in individual commodity prices. We find that our DFM constitutes a crude, but effective, tool for analyzing commodity price movements. Studying Atlantic salmon and Brent Crude, we find that a moderate, but significant, fraction of their price movements over the sample period (1980-2018) can be attributed to shocks to a global factor representing common, demand-driven trends in global commodity markets. Notably, a sub-sample analysis reveals that since 2000, the global factor's importance has increased significantly for both commodities, indicating an increased level of integration with global commodity markets.

In the second study, we use the DFM framework to forecast the prices of Atlantic salmon and Brent Crude oil. Specifically, the monthly spot price is predicted 1-3 months ahead using data from 2006 to 2018. We assemble a comprehensive set of predictors for each commodity, and carry out model selection by employing a genetic algorithm. The resulting models' out-of-sample performances are assessed and compared against more commonly used forecasting models and results from the literature. We find that our forecasts improve upon all benchmarks across all horizons for both commodities. Our results indicate that there is value to be gained from forecasting based on *latent*, estimated factors representing co-variation *within* a set of recognized predictor variables, rather than based on the predictor variables directly.

Sammendrag

Formålet med denne oppgaven er å studere prisbevegelsene til to viktige norske eksportvarer, atlantisk laks og Brent råolje, gjennom en tilærming basert på dynamisk faktormodellering (DFM). Vi utfører to ulike analyser, en sambevegelse- og komposisjonsanalyse av prisendringer i råvarer, samt predikering av månedlige spotpriser.

I den første analysen utforsker vi sambevegelser i globale råvaremarkeder gjennom å anvende en dynamisk faktormodell til å dekomponere avkastningen på råvarepriser inn i globale, sektor-spesifikke og idiosynkratiske komponenter. Denne dekomponeringen tillater oss å estimere til hvilken grad globale, sektor-spesifikke og råvare-spesifikke faktorer og hendelser forklarer hver enkelt råvares prissvingninger.

I den andre analysen bruker vi DFM-rammeverket til å predikere prisene på atlantisk laks og Brent råolje, nærmere bestemt den månedlige spotprisen 1-3 måneder frem, ved å bruke data som spenner fra 2006 til 2018. Vi konstruerer et omfattende sett med prediktorer for hver råvare, og utfører modellseleksjon ved å benytte en genetisk algoritme. De resulterende modellenes prediksjonsevne blir vurdert og sammenlignet med ofte anvendte predikeringsmodeller, samt resultater fra litteraturen. Vi finner at våre resultater slår alle benchmarks over alle tidshorisonter for både atlantisk laks og Brent råolje. Disse resultatene indikerer at det ligger verdi i predikering basert på *latente*, estimerte faktorer som representerer samvariasjon innad i et sett med anerkjente prediktorvariabler, i stedet for basert på prediktorvariablene direkte.

Table of Contents

List of Figures

List of Tables

Abbreviations

ADF Augmented Dicky Fuller

AIC Akaike Information Criterion

AR Autoregression

ARCH Autoregressive Conditional Heteroscedasticity

ARIMA Autoregressive Integrated Moving Average

BDI Baltic Dry Index

BCO Brent Crude oil

BVAR Bayesian Vector Autoregression

CAD Classical Additive Decomposition

DFM Dynamic Factor Model

D-M Diebold-Mariano (1995)

EM Expectation Maximization

E&P Exploration and Production

E-step Expectation Step

GA Genetic Algorithm

GFC Global Financial Crisis

GWE Gutted Weight Equivalents

JB Jarque-Bera

MASE Mean Absolute Scaled Error

ML Maximum Likelihood

MLE Maximum Likelihood Estimate

MoM Month-on-Month

MSPE Mean Squared Prediction Error

M-step Maximization Step

OLS Ordinary Least Square

OECD The Organization of Economic Co-operation and Development

OPEC The Organization of Oil Exporting Countries

PANIC Panel Analysis of Non-stationary Idiosyncratic Compoents

PC Principal Components

PSALM Norwegian Export Price of Farmed Atlantic Salmon

RMSE Root Mean Square Error

rRMSE Relative Root Mean Square Error

SSB Statistics Norway

VAR Vector Autoregression

WTI Western Texas Intermediate

YoY Year-on-Year

1 Introduction

Insights into the sources/drivers, co-movement and future paths of commodity prices are valuable to a wide range of stakeholders. Among these are a nation's economic institutions, due to movements in commodity prices feeding through to real economic activity and thus influencing fiscal and monetary policy considerations. It is of particular importance for a small, open, commodity-dependent economy like that of Norway to understand the price dynamics of its major commodities. Further, insights into the future price trajectory of a commodity is of obvious interest to industry participants along all parts of the value chain. Commodity producers make their production or extraction decisions on the basis of what they believe the future price path will be. Similarly, value-added providers continuously plan and adjust their operations according to the expected costs of their main input factors, i.e. the prices of commodities. Financially, an understanding of price behaviour and access to accurate price predictions are essential to risk management, stock and bond valuations, as well as a multitude of investment decisions.

However, gaining genuine insights into the dynamics of commodity prices is a complex task. Prices are affected by a wide variety of factors, ranging from environmental and biological factors, to more fundamental economic factors. Furthermore, over the last decades, commodity markets have become subject to increased financialization - leading to further increased complexity of price dynamics (Cheng and Xiong, 2014). The complexity of commodity markets is illustrated in the World Bank's note on the food price surge from 2006 to 2008, concluding that drivers of price change included factors such as long-term demand and supply, higher energy prices, increased bio-fuel production, depreciation of the U.S. dollar, adverse weather conditions and monetary policy responses (Mitchell, 2008). Another example can be found with the price of oil. According to Hamilton (2008), the oil price's statistical properties largely resemble that of a random walk and is therefore an economic variable that should be close to impossible to predict. In short, valuable insights into the movements of commodity prices are inherently difficult to obtain.

This fact notwithstanding, the aim of our thesis is to investigate the price movements of two commodities highly material to the Norwegian economy: Atlantic salmon and Brent Crude oil. With this selection, we examine one commodity from each of what are arguably the two most interesting sectors in a Norwegian context, namely the offshore and seafood sectors. The offshore sector is the country's largest export sector by far, responsible for about 50% of Norwegian export value (SSB, 2018). Seafood is the largest product category in the country's second most important export sector (main land), and considering the aquaculture industry's growth ambitions along with the expected future decline in petroleum-related revenue, its importance seems destined only to increase.

In order to best analyze and model our selected commodities and capture their complex time series dynamics, we have chosen a dynamic factor modelling approach. Dynamic factor models (DFMs) efficiently summarize the information contained in large data sets by extracting a small number of latent factors representing the co-variation within the set. Common to factor models in general is the idea that the observed variables can be modelled as a linear combination of factors. What separates DFMs from other factor models is the fact that their factors are unobservable - they are extracted from the given data and capture underlying trends. Further, time dynamics are incorporated, which allows for both lagged and contemporaneous relationships. These features enable modelling of complex relationships among large sets of variables. In any case, the main advantage of a DFM approach is enhanced information utilization and dynamic modelling abilities relative to standard time-series analysis[1], without the loss of parsimony.

DFMs are increasingly used in data rich environments. The general model framework originates from Sargent and Sims (1977) and Geweke (1977). Their motivation was based on the notion that if a variable is affected by a large number of different factors, then *all* these factors contain useful information and should be included in the analysis and modelling of this variable. The model framework combines cross-sectional and standard time-series analysis, which has made it compelling to the fields of economics and finance where common shocks can drive co-movement of a large number of interrelated variables. As a result, economists are increasingly looking to these models for policy analysis (Forni et al., 2005). Recent applications of the framework includes investigating the world's risky asset markets, factor analysis and now-casting of gross domestic products (GDP), and forecasting of housing prices, to mention a few (e.g., see Miranda-Agrippino and Rey, 2015; Giannone et al., 2008; Emiris, 2016).

[1]Such as linear regression and vector autoregressive (VAR) models.

As can be seen, representing economic variables through a dynamic factor model set-up opens many interesting avenues of research. In this study we focus on two specific commodity-related applications, namely *co-movement and composition analysis* of commodity returns and *forecasting* of commodity spot prices. Henceforth, we refer to these two studies as *Application A* and *Application B*, respectively.

In *Application A* we investigate dynamic co-movement in global commodity markets via a DFM applied to a large data set of globally traded commodities. This approach enables us to decompose commodity price returns into global, sectoral, and commodity-specific components. Specifically, we estimate the share of individual commodities' price movements attributable to each of the following sources: 1) broader, common trends in global commodity markets, 2) common trends within sub-groups of related commodities, and 3) shocks unique and isolated to the individual commodity markets. In our analyses, an emphasis is put on the price movements of Atlantic salmon and Brent Crude oil. The insights obtainable from this application of the DFM should be of interest to actors ranging from salmon farmers to the Norwegian Central Bank.

In *Application B*, we generate 1-3 month-ahead forecasts for Atlantic salmon and Brent Crude oil spot prices by combining the use of DFMs with a comprehensive set of predictors. We specify the DFMs optimally by using a genetic algorithm for model selection. In the last decades, dynamic factor modelling has become a popular tool for forecasting macroeconomic variables, for which it produces solid results (e.g., see Bańbura and Modugno, 2014; Emiris, 2016; Jansen et al., 2016). With Application B, we break new ground by testing whether the approach can be successfully applied to forecasting commodity prices as well, and thus provide potentially significant economic value to a wide range of stakeholders.

Each of the applications described require their own unique adaptation and implementation of the general DFM framework. The two applications also provide distinctly different insights. So why include two such different applications in the same study? Because together they constitute both a comprehensive analysis of the commodity markets we seek to investigate, as well as a solid testament to the general DFM framework's usefulness and applicability to the modelling of commodity prices.

Our thesis is structured in the following manner. In Section 2, we discuss some characteristics of the markets/industries for both Atlantic salmon and Brent Crude oil. Insights from this section will inform the analyses performed in the subsequent sections. For readers somewhat unfamiliar with these commodity markets, the section should provide them with the necessary fundamentals.

In Section 3, we present the general econometric model framework, and show how a dynamic factor model cast in state-space form constitutes the basis for the work in both Application A and Application B. In this section we also explain how we solve the complex task of estimating the dynamic factor models through maximum likelihood estimation, which we implement using an Expectation-Maximization (EM) algorithm combined with a Kalman smoother, and how the employment of such algorithms offer several benefits in empirical modelling.

Section 4 is dedicated to Application A. First, we provide motivation for, and discuss literature relevant to, this specific application. Our data set of globally traded commodities retrieved from the World Bank is subsequently presented, before our specific methodology is laid out. We explain how we through constructing and estimating a *restricted* dynamic factor model are able to obtain a three-level factor representation of our commodity time series, and assign economic interpretations to each of these factors. Finally, we present and discuss our empirical results. We succeed in providing novel insights into the extent to which global, sectoral, and commodity-specific shocks contribute to the price movements of Atlantic salmon and Brent Crude oil.

Section 5 is dedicated to Application B. First, we provide motivation for this application of the DFM framework and review relevant literature. Second, we introduce our modelling approach; we present the set of predictor variables, the benchmark models, and explain how we perform model selection using a genetic algorithm. Finally, we assess our forecasts and compare their performance against benchmark models and results found in comparable studies. Our DFM approach produces promising results, as we outperform the benchmark models across all horizons for both commodities.

Section 6 wraps up our thesis. Here we summarize the findings from Application A and Application B, provide concluding remarks, as well as suggestions to further research and how our work can be improved and extended upon.

2 Industry and Market Characteristics

Brief introductions to the markets and industries of what are arguable two of the most interesting commodities in a Norwegian context, Atlantic salmon and Brent Crude oil, are provided below.

2.1 Atlantic Salmon

World population is expected to exceed nine billion in 2050, putting enormous pressure on our ability to produce sufficient amounts of food. Resources for increased land based protein production is likely to be scarce, and significant growth in capture fisheries on already over-exploited wild fish stocks seems unlikely. Many therefore look to the aquaculture industry to fill the coming supply-demand gap (Marine Harvest, 2017). And today, aquaculture is in fact the fastest growing food production system in the world, and as of 2014 aquaculture production constitutes more than half of all fish consumed by humans (FAO, 2016). Simply put, farmed fish is a commodity for the future.

Atlantic salmon is a leading species within today's modern aquaculture industry.[2] The most significant producer countries of this species are Norway, Chile, Canada, Great Britain and Faroe Islands (Solibakke, 2012). Out of these, Norway and Chile are the most important producers, with respectively 54 and 23 percent of the market share in terms of volume. Total production in 2016 was around 2 million tonnes gutted weight equivalents (GWE) (Marine Harvest, 2017).

In the past, the aquaculture industry was characterized by many small businesses, but has gradually consolidated through mergers and acquisitions. Vertical integration has also been increasing in the industry, with large companies now holding ownership in the entire value chain, from smolt production to processing plants. The largest Norwegian salmon aquaculture companies today are Marine Harvest, Salmar and Leroy Seafood (Marine Harvest, 2017).

[2]Atlantic salmon's leading position is due to decades of industrialization and increasing productivity, where technological development and productivity growth has gradually lead to lower production costs, increased competitiveness, and lower prices for consumers (Asche et al., 2013). Though considered a leading species, Atlantic salmon production still constitutes only a small share of global aquaculture in terms of volume produced, and there exists a huge potential for industrialization in the rest of the aquaculture sector.

Supply and Demand Characteristics

The Atlantic salmon industry is characterized by tight supply-demand conditions, and shifts in both factors display strong explanatory power with regards to Atlantic salmon price dynamics (Oglend, 2013). It is also characterized by a long production cycle and strong seasonality effects. Due to limited opportunity to adjust production volumes in the short-to-medium term given the long production cycle, positive (negative) demand effects are to a large extent adjusted by positive (negative) price movements rather than increased (decreased) supply. This indicates a potential for strong demand-side effects on price movements from consumption markets. Brækkan and Thyholdt (2014) measured demand shifts[3] in all salmon-importing regions of the world using an index approach for the 2002-2011 period. They found that demand growth for salmon is highly unstable and characterized by large variations between regions and over time within regions.[4] It is indicated that these variations partly explain the high volatility of salmon prices over the last decade.

At the same time, a range of studies have found tight correlations between shifts in supply-side factors, such as costs, and salmon price movements (e.g., see Oglend, 2013; Asche and Bjorndal, 2011). As the industry has matured, costs have shifted from being productivity driven, to becoming more affected by input-factor prices (Asche and Oglend, 2016). Prices tend to be higher when costs increase, indicating the presence of pass-through effects.

In the last decade, the industry has seen stagnation in productivity and increased costs due to higher feed prices, as well as biological issues such as parasites and disease (Marine Harvest, 2017). The recent surge in salmon prices we see in Figure (1) is considered to be a product of strong demand growth, supply-limitations due to disease problems and stringent regulations, as well as upwards-trending production costs. The weighting and relative importance of these factors however, is not entirely clear.

[3] Shifts in quantity demanded for a given price.

[4] We note that total global demand growth for salmon in this period was about 94 percent.

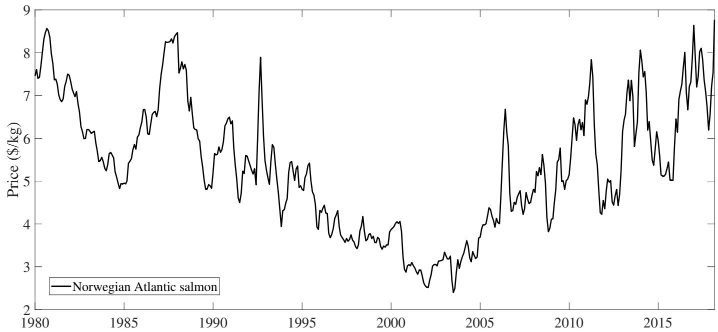

Figure 1: *The historical development of the price of Atlantic salmon. The price of farmed Atlantic salmon has increased substantially in recent years. This new trend follows a long period of consistently falling prices. At the same time, price volatility has increased.*

2.2 Crude Oil

Oil is a global commodity that is sold and delivered to most corners of the world. After crude oil is extracted from a field, it is sent to a refinery where the various hydrocarbons are separated, transformed and processed in order to make oil products. Oil products are used mainly as fuel, with the transport sector being the largest consumer. Covering about 32 percent of the world's total energy demand, it is the world's largest source of energy, followed by coal and gas (International Energy Agency, 2017). However, an increasing proportion is now used outside the energy sector, with the petrochemical industry being the second largest consumer of oil (International Energy Agency, 2017).

Saudi Arabia was the world's largest producer of crude oil in 2016, followed by Russia and the United States (International Energy Agency, 2017). Norway is a relatively small player in the global oil market, with its production covering about two percent of the world's global demand for crude oil. However, almost all oil produced on the Norwegian shelf is exported, and accounts for about 25% of total Norwegian export value (SSB, 2018).

The international oil industry consists of a multitude of different players involved in various stages of the value chain, including exploration & production (E&P), transport, processing and marketing. The major E&P companies can be divided into the following two groups: State-owned companies and international oil companies. State-owned companies in the producer country often play a significant financial and administrative role, and many are also technically and operationally active. Examples of state-owned companies are Equinor, Saudi Aramco, National Iranian Oil Co., and Pemex. Among the international oil companies, the most important actors are considered to be Shell, Exxon-Mobil, BP and Chevron (including Texaco).[5]

Although largely treated as a singular commodity in this section, it is important to note that there exists a multitude of types. The crude oil types differ in desirability, determined by characteristics such as density, sulphur level, and extraction location. Buyers of crude oil need an easy way to value the commodity based on its quality and location. Benchmark oils such as Brent, Western Texas Intermediate (WTI) and Dubai/Oman serve this purpose. The Brent benchmark is the most widely used marker, referenced in about 2/3 of crude contracts (Wall Street Journal, 2018). It refers to oil from four different fields in the North Sea. Crude from this region is light and sweet, and thus ideal for the refining of diesel, petrol and other strongly demanded products. Being extracted offshore, it is also easy to transport to distant locations. The historical price series of Brent is depicted in Figure (2).

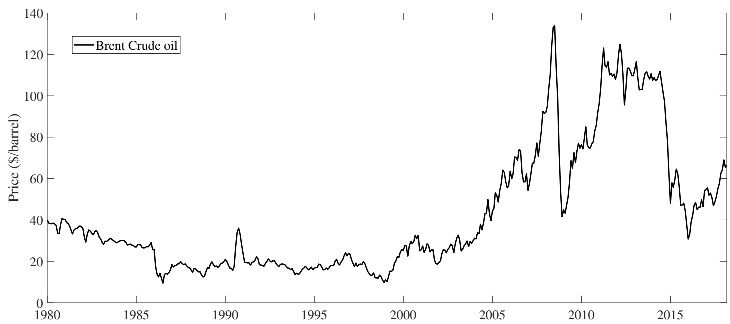

Figure 2: *Historical development for the spot price of Brent Crude oil.*

Supply and Demand Characteristics

Crude oil is traded on a global scale among many different actors: oil producing countries, oil importing countries, oil companies, refineries, and speculators. In the long run, oil is considered a fairly elastic commodity, where movements on the supply and demand, or production and consumption, sides are well reflected in the price (Hagen, 1994; Stevens, 1995).

[5] These are generally integrated companies, i.e. involved in all stages of the value chain.

Oil consumption is naturally linked to global macroeconomic aggregates such as real GDP growth and Industrial Production. Demand for oil, both consumer market and industrial, tends to increase (decrease) along with increasing (declining) economic growth rates. Unexpected changes in oil consumption driven by the global business cycle has been shown to yield high explanatory power for various price shocks, such as the 2003-2008 surge (Kilian, 2009).

On the supply side, a known characteristic of the oil industry is that market control has traditionally been concentrated on a few actors. Initially, it was primarily the major international oil companies that had dominant influence through cartels and cooperation agreements. These companies still have a very strong position, but over the years, national governments have gained a stronger control over production through their national companies. This change has been of great significance, with oil prices to a larger extent being affected by political considerations. The Organization of the Petroleum Exporting Countries (OPEC), operating as a cartel with the largest production capacity, has served a decisive role in the price formation of crude oil since the 1960s (Fattouh, 2007).[6] However, recent developments, such as the US shale oil revolution and reduced OPEC spare capacity, have contributed to limiting the organization's control on global production levels, abating their function as the major buffer on the supply side. OPEC's diminishing ability to operate as swing producer has been identified as an important factor for the price fluctuations observed in recent years (Olimb and Odegaard, 2010).

Finally, non-fundamental factors such as speculative demand and developments in the futures markets, have also been shown to significantly impact oil price determination (e.g., see Kilian and Murphy, 2014; Olimb and Odegaard, 2010; Westgaard et al., 2017).

The interplay of the factors discussed, along with others left unnoted in this brief review, have lead to a strongly fluctuating crude oil market, providing it with characteristics such as complex non-linearity, dynamic variation and high irregularity (Plourde and Watkins, 1994).

[6] Oil production from OPEC comprised about 40 percent of the world's oil supply in 2016. In addition, OPEC holds 3/4 of total reserves (International Energy Agency, 2017).

3 Econometric Framework

The econometric framework used throughout this thesis is a dynamic factor model estimated via the Expectation-Maximization (EM) algorithm. The adopted framework is based on the comprehensive work of Doz et al. (2012) and Bańbura and Modugno (2014). Next, we review the literature pertaining to this econometric framework and state the rationale and motivation for using it. In Section 3.1 we present the general dynamic factor model. In Section 3.2 we describe the estimation procedure, more specifically the EM algorithm.

The steadily increasing volume of financial and economic data available has spurred the need for methods that are able to efficiently exploit the information obtainable from large data sets. Several methods have been proposed for empirical modelling in such data-rich environments. One popular approach is the *general-to-specific* approach, in which a general model that adequately fits the evidence is simplified, e.g. through variable reduction (e.g., see Hoover and Perez, 1999; Krolzig and Hendry, 2001). This approach is limited by the fact that one inevitably has to leave out potentially important variables. More conventional models, such as regressions and vector autoregressive (VAR) models, suffer from over-parameterization when the number of variables become large. To remedy this problem, Bayesian estimation methods introduce the notion of informative priors which opens the possibility of imposing informed restrictions on these models, hence reducing dimensionality. As a result, Bayesian VAR (BVAR) models have become popular in forecasting (e.g., Karlsson, 2013; De Mol et al., 2008).

Dynamic factor modelling originated with the work of Sargent and Sims (1977) and Geweke (1977), and has increasingly gained attention as a method for empirical modelling in data-rich environments. Different from other large-scale models, the DFMs represent a flexible framework that handles over-parameterization problems by summarizing large-scale information into a small number of latent factors. More specifically, each observed variable is modelled as a sum of factors, where each factor represents underlying trends in the data. By doing so, noise is reduced and genuine relationships are captured in a parsimonious manner. Such appealing properties have made DFMs extensively used in econometric applications (e.g., see Forni et al., 2000, 2005; Giannone et al., 2008; Bańbura et al., 2011; Miranda-Agrippino and Rey, 2015). For a comprehensive literature study on dynamic factor models, we refer to Barhoumi et al. (2013).

There are several methods for estimating the latent factors in a DFM. For a long time, principal components (PC) were used as factor estimates in lack of any better estimators (Bai, 2003; Forni et al., 2009). Doz et al. (2011, 2012) have later shown that maximum likelihood estimation (MLE) is suitable for DFMs, and that it is robust to model misspecification.[7] Maximum likelihood estimation can be implemented using the Kalman filter/smoother and the Expectation-Maximization algorithm. The EM algorithm was first proposed by Dempster et al. (1977) and Shumway and Stoffer (1982) as an iterative algorithm to find maximum likelihood estimates of parameters in statistical models where the model depends on unobserved latent variables. The EM algorithm solves the problem related to situations where optimization of the likelihood function is analytically intractable, as it is in the case of DFMs, due to the latent variables and unknown model parameters that must be estimated simultaneously. A large body of research has made use of the fact that DFMs may be estimated through maximum likelihood estimation (e.g., see Delle Chiaie et al., 2017; Emiris, 2016; Miranda-Agrippino and Rey, 2015).

The rationale for using a state-space model (such as a DFM) over more conventional methodologies such as ARIMA, VAR or other ad hoc/heuristic methods, is manifold. Although the "ARIMA family" might be considered advantageous due to being a universal estimator[8] with ease of both implementation and estimation, state space models allow for modelling of more complex processes of large dimensions. In general, our model allows for; (i) dynamic model parameters that automatically handle and adjust for structural breaks, shifts and time-varying dynamics; (ii) the use of different time series simultaneously to estimate one underlying quantity[9]; (iii) the construction of interpretable model structures embedded in economic theory; (iv) modelling of underlying drivers as VAR processes, including lead/lag relations into the model. These traits make the DFM ideally suited for our dual-purpose thesis.

[7]More specifically, it is robust to cross-sectional and time series correlation of the idiosyncratic components. The implications of this will be further discussed.

[8]The Wold decomposition theorem states that any co-variance stationary process can be decomposed into two mutually uncorrelated component processes; (i) a linear combination of lags of a white noise process and (ii) a process which future values can be predicted exactly by a linear function of past observations. This property makes ARIMA models popular for forecasting purposes (Brooks, 2014).

[9]By this we refer to different time window and time series character (units and properties).

3.1 The Dynamic Factor Model

Let $y_t = [y_{1,t}, y_{2,t}, \ldots, y_{n,t}]'$ be an n-dimensional vector of time series that all satisfy the assumption of stationarity. We assume that y_t admits the following factor model representation:

$$y_t = \Lambda F_t + \epsilon_t, \tag{1}$$

where $F_t = [f_{1,t}, ..., f_{r,t}]'$ is an $r \times 1$ vector of (latent) common factors. The variables y_t load on the factors through the factor loadings in Λ, which is an $n \times r$ matrix. ϵ_t is a vector of idiosyncratic components which capture the remaining variation not explained by the common factors. The common factors and the idiosyncratic components are assumed to be zero-mean, stationary processes.

Further, the factors are modelled as a VAR process of order p:

$$F_t = A_1 F_{t-1} + ... + A_p F_{t-p} + u_t \qquad u_t \sim i.i.d\ N(0, Q) \tag{2}$$

The autoregressive coefficients are collected in p matrices $A_1, \ldots, A_p$, each of size $r \times r$. In contrast to static factor models, dynamic factor models provide a rich representation of the data and exploits potentially crucial information contained in the lead and lag relationships between the factors (Forni et al., 2005). Taking the factor dynamics explicitly into account is particularly important in forecasting applications.

Finally, we assume that the idiosyncratic components follow AR(1) processes:

$$\epsilon_t = \rho \epsilon_{t-1} + e_t \qquad e_t \sim i.i.d\ N(0, R) \tag{3}$$

where ρ is a diagonal matrix of size $n \times n$. The common factors and the idiosyncratic components are assumed to be uncorrelated.

The DFM is said to be exact if the idiosyncratic components (ϵ_t) are both serially and cross-sectionally uncorrelated. In an exact model, all the dynamic interactions between the observable variables can be attributed to the r common factors (Emiris, 2016). However, the assumptions of an exact DFM are often not satisfied by the data. The common factors are often not able to capture all the co-variation existing in large data sets if the variables are dissimilar in nature. Fortunately, Doz et al. (2012) have shown that such approximate DFMs can be estimated by maximum likelihood under the assumption of serially and cross-sectionally uncorrelated idiosyncratic components even if this condition is not satisfied by the data.

Even though the model can be viably estimated when these assumptions are mildly violated, we aim to further reduce the degree of misspecification. First, we allow for several types of common factors. Global factors will represent underlying trends common to all the variables - all variables will load on these factors. Block factors will represent underlying trends common only to a subset of variables - only the variables in a given subset will be able to load on the corresponding factor. This decreases the likelihood of correlated idiosyncratic factors, as we can create groups of variables likely to have co-varying relationships. In the most general form of the model, the factor representation consists of any number of global factors and any number of block factors. Further, as shown above, we choose to model the idiosyncratic components (ϵ_t) as AR(1) processes. Modelling serial correlation explicitly should leave the residuals e_t at most weakly correlated across variables.

Equations (4) and (5) make up the matrix representation of the dynamic factor model setup just described.

$$\begin{pmatrix} y_{1,t} \\ y_{2,t} \\ \vdots \\ y_{n,t} \end{pmatrix} = \begin{pmatrix} \lambda_{1,1} & \lambda_{1,2} & \cdots & \lambda_{1,r} \\ \lambda_{2,1} & \lambda_{2,2} & & \vdots \\ \vdots & & \ddots & \lambda_{n-1,r} \\ \lambda_{n,1} & \cdots & \lambda_{n,r-1} & \lambda_{n,r} \end{pmatrix} \begin{pmatrix} f_{1,t} \\ f_{2,t} \\ \vdots \\ f_{r,t} \end{pmatrix} + \begin{pmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \\ \vdots \\ \epsilon_{r,t} \end{pmatrix} \tag{4}$$

$$\begin{aligned} \begin{pmatrix} f_{1,t} \\ f_{2,t} \\ \vdots \\ f_{r,t} \end{pmatrix} &= \begin{pmatrix} a^1_{1,1} & a^1_{1,2} & \cdots & a^1_{1,r} \\ a^1_{2,1} & a^1_{2,2} & & \vdots \\ \vdots & & \ddots & a^1_{r-1,r} \\ a^1_{r,1} & \cdots & a^1_{r,r-1} & a^1_{r,r} \end{pmatrix} \begin{pmatrix} f_{1,t-1} \\ f_{2,t-1} \\ \vdots \\ f_{r,t-1} \end{pmatrix} + \cdots \\ &+ \begin{pmatrix} a^p_{1,1} & a^p_{1,2} & \cdots & a^p_{1,r} \\ a^p_{2,1} & a^p_{2,2} & & \vdots \\ \vdots & & \ddots & a^p_{r-1,r} \\ a^p_{r,1} & \cdots & a^p_{r,r-1} & a^p_{r,r} \end{pmatrix} \begin{pmatrix} f_{1,t-p} \\ f_{2,t-p} \\ \vdots \\ f_{r,t-p} \end{pmatrix} + \begin{pmatrix} u_{1,t} \\ u_{2,t} \\ \vdots \\ u_{r,t} \end{pmatrix} \end{aligned} \tag{5}$$

In the upcoming empirical applications, we will be interested in restricting the model parameters Λ and $A_1, \ldots, A_p$ in order to tailor the DFM to the given application. Through imposing restrictions, we are able to create model structures based on economic theory, allowing us to interpret the model in a more structural sense. However, imposing such restrictions comes at the cost of increased complexity in the implementation of the estimation procedure.

In order to estimate the DFM, it must be cast in *state space* form. This requires us to separate the different elements of the model, namely the various factors in the model, collectively referred to as the *state vector*, and the *model parameters*. For a general model with n variables represented by r factors following VAR(p) processes, such as the model stated in Equations (1)-(3), the following state vector x_t and model parameters θ need to be estimated:

$$\begin{aligned} x_t &= (f_{1,t}, \ldots, f_{r,t}, \epsilon_{1,t}, \ldots, \epsilon_{n,t}), \\ \theta &= (\Lambda, A_1, \ldots, A_p, Q, \rho_1, \ldots, \rho_n, R) \end{aligned} \tag{6}$$

The actual state space representation is more of a technical matter, and we refer to Appendix E for the full state space representation of the models used in *Application A* and *Application B*.

3.2 Estimation

We choose to estimate the DFM by maximum likelihood estimation (MLE), which we implement using an EM algorithm combined with a Kalman smoother. This is a viable estimation method for DFMs, as shown by Doz et al. (2011, 2012). EM algorithms extend maximum likelihood estimation to models with hidden states, such as the DFMs' latent factors. While the EM algorithm provides the framework of the estimation procedure, the Kalman smoother is employed to actually produce the latent factor estimates. It is an extension of the more well-known Kalman filter, which is a widely applied concept in time series analysis and econometrics where it is used to estimate unknown variables from observable, noisy time series. For the interested reader, we have produced a full technical description of the estimation procedure and its elements, which can be found in Appendix F.

Next, we introduce the main concepts behind the estimation procedure, as this should give a sufficient understanding of how the DFMs are estimated throughout this thesis. In short, the EM algorithm seeks to find the maximum likelihood estimates by iteratively improving the estimates of the state vector x_t and model parameters θ. This is done in two separate steps: the Expectation step and the Maximization step. The key elements of the full EM algorithm are:

- **Initialization.** The algorithm is initiated by finding estimates of both the state vector (x_t) and model parameters (θ). The initial model parameters are estimated by performing an Ordinary Least Square (OLS) regression on the principal components of the observable variables, treating these as the factors. The initial state vector is estimated via the Kalman filter/smoother, taking the model parameter estimates as given.

After initialization, the algorithm iterates between the expectation step (E-step) and the maximization step (M-step). We denote the current iteration by k.

- **E-step.** The *Kalman filter/smoother* is used to estimate the expected value of the state vector, x_t, given the current estimate of model parameters and the observed data; that is, $x_t^k = \mathbb{E}_{\theta^{k-1}}\big[x_t | y_1, \dots, y_t\big]$.
- **M-step.** The model parameters are estimated by finding the parameters that maximize the log likelihood function, taking the state vector as given. That is, $\theta^k = \text{argmax}_\theta \, \mathbb{L}\big[x_t^k, \theta\big]$, where $\mathbb{L}$ is the log likelihood function. Technically, this is done by differentiating the log likelihood function, setting it equal to 0, and solving for the model parameters.

The E-step and M-step is repeated until a defined convergence criterion is met. The convention is to compute the following convergence criterion after the M-step;

$$c_k = \frac{\mathbb{L}\big[x_k, \theta_k\big] - \mathbb{L}\big[x_{k-1}, \theta_{k-1}\big]}{\frac{1}{2}\big(\mathbb{L}\big[x_k, \theta_k\big] + \mathbb{L}\big[x_{k-1}, \theta_{k-1}\big]\big)}, \qquad c_k \overset{?}{<} \tau, \tag{7}$$

where τ is the threshold that, when reached, terminates the estimation procedure. We stop after iteration k if $c_k < \tau$.

In order to satisfy the model requirements of our studies, we need to make certain modifications to the general model framework presented by Doz et al. (2012). First, we need to be able to model time series of varying lengths and incomplete time series in general. Bańbura and Modugno (2014) show how the maximum likelihood estimation of dynamic factor models can be done when dealing with arbitrary patterns of missing data, and so we implement their proposed changes. Second, we need to be able to impose restrictions on the model parameters. In both applications we restrict the loading matrix Λ, and in Application A we also restrict the model parameters Q, and $A_1, \ldots, A_p$. We combine insights from Wu et al. (1996), Bork (2009), and Bańbura and Modugno (2014) so that the EM algorithm accommodates these restrictions. In technical terms, this is done through placing linear restrictions of the form $H_\Lambda \operatorname{vec}(\Lambda) = \kappa_\Lambda$ and $G_A \operatorname{vec}(A) = \delta_A$, where the matrices H and G are selection matrices that select the elements of Λ and A that are to be restricted. The vectors κ and δ contain the constants of the linear restrictions. The full estimation procedure, including the Kalman filter/smoother and the final maximization equations of the model parameters, can be found in Appendix F.

4 Application A: Co-Movement and Composition Analysis

Recall that in this application, we seek to investigate dynamic co-movement in commodity markets via a dynamic factor model approach and decompose commodity returns into global, sectoral, and idiosyncratic components. This decomposition will allow us to estimate the degree to which global, sectoral and commodity-specific factors contribute to the movements in individual commodity returns. Our main subjects of analysis are Atlantic salmon and Brent Crude oil.

Increased understanding of co-movement between commodity markets, as well as the sources of changes to individual commodity prices, is valuable to a wide range of actors. For a nation's economic institutions, knowing the extent to which specific price movements in its imports and exports are driven by global, sector-specific or commodity-specific shocks is highly valuable as this affects the impact these movements may have on a country's real economy. Illustrating this, Bjørnland and Thorsrud (2014) studied how changes in the oil price affect Norwegian mainland GDP growth, depending on whether it is caused by global demand shocks or by a shock in the global supply of oil. They found that an oil price decrease caused by a global demand drop leads to a considerably greater decrease in domestic GDP compared to a price decrease due to over-supply. Kilian (2009) reports similar findings when studying the effect of oil price shocks on the US economy. Other major benefits to understanding the composition of commodity price movements, relate to macroeconomic concepts such as the Dutch disease[10] and imported inflation. For a deeper understanding of the usefulness of this form of insight, we refer to the works of e.g. Bjørnland (1997), Bjørnland and Thorsrud (2016) and Tang et al. (2014). For industry and financial players, insights into individual commodities' integration with global markets can inform e.g. risk management, hedging strategies and investment decisions.

[10]The effects of higher prices of natural resources on an exporting country's real exchange rate: an appreciation likely leading to substantial reallocations of factors of production, and a decline in exports of other goods. Price surges due to common trends in global commodity markets are naturally less likely to contribute to this phenomena.

Section 4 is organized in the following manner. First, we review literature relevant to this application. Second, we introduce the information set of globally traded commodities utilized in our study, before we outline our methodology. Lastly, we discuss our findings/empirical results.

4.1 Literature Review

Co-movement in commodity markets has been thoroughly studied over the years, and we now highlight some of the work most relevant to our study. Cheng and Xiong (2014) find that commodity price dynamics have substantially changed over the past decade. They emphasize the following; First, even across sectors, commodity prices have tended to move together as a class since the early 2000s; Second, correlation between commodities and other asset classes is strictly positive, and increasing, in the same period.[11] Further, Kilian (2009) investigated the oil price shock of 2003 to 2008, and found that the surge in the real price of oil was driven by repeated positive demand for *all* industrial commodities. Juvenal and Petrella (2015) find that co-movement among energy and non-energy commodity prices increases due to global demand shocks. These authors point at the unexpected high growth in emerging markets, mainly in China and other emerging economies in Asia, as the main reason for these observations. Delle Chiaie et al. (2017) argue that global demand shocks from global economic activity are likely to affect most commodities by shifting the demand curve in the same manner across markets. Large amounts of research have been devoted to investigating co-movement in commodity markets. The consensus from the literature referred to above suggests that the increasingly integrated markets of commodities respond similarly to shocks related to economic activity. This insight is exploited in our application of the DFM framework.

We are not the first to use a factor approach to seek insights into commodity markets (e.g., see Alquist and Coibion, 2014; Byrne et al., 2013; Delle Chiaie et al., 2017; West and Wong, 2014; Yin and Han, 2015). Common to all these studies is the utilization of large data sets of traded commodities, for the purpose of examining co-movement and price drivers present in commodity markets. Alquist and Coibion (2014) create a theoretical factor structure for commodity prices where the common factor captures the combined contribution of all aggregate shocks

[11] Assets classes such as equity, exchange rates and US treasury bills. They state that high correlation with equities may be a result of increased financialization of commodity markets.

that affect global commodity markets. Using a Panel Analysis of Nonstationary and Idiosyncratic Components (PANIC) approach, Byrne et al. (2013) document a significant degree of co-movement between globally traded commodities. With their factor model, West and Wong (2014) find that individual commodity prices show a tendency to revert towards one common factor. The approaches perhaps most similar to our own, are Delle Chiaie et al. (2017) and Yin and Han (2015), who both use a three-level dynamic factor model to decompose commodity returns into global, sectoral, and commodity-specific components. Interestingly, Delle Chiaie et al. (2017) indicate that the decomposition provided by their DFM can be interpreted along a supply/demand dimension, where the global factor mainly captures global commodity demand, while the idiosyncratic factors mainly capture industry-specific supply shocks.

This study in context of the existing literature

Our approach is inspired by the body of work discussed above. That being said, the work we perform here is differentiated from the existing literature on factor modelling of commodities in several ways. First, most of these papers studied commodity prices in levels instead of returns, focusing primarily on co-integration relationships and long term trends. Along with Delle Chiaie et al. (2017) and Yin and Han (2015), we focus on price movements in terms of monthly log returns. Of the studies listed, these are also the only ones applying a *dynamic* factor approach. From these two papers, we differ in other ways. *First*, we use a different, and more updated, set of world commodity data, as well as the employment of a different model structure. *Second*, where these works mainly focus on general findings regarding co-movement in global commodity markets, we narrow in on commodities material to the Norwegian economy. Our study is the first to perform a detailed analysis of Atlantic salmon and Brent Crude oil, and their co-movement with global commodity markets, using a DFM approach.

4.2 Data

To perform our analysis, we use a monthly data set of globally traded commodities spanning the time period 1980-2018. The data set is retrieved from the World Bank's commodity database,[12] which represents the diversity of the world's major commodity markets. As Atlantic salmon is not included in the World Bank's set of commodities, we retrieve its price series from Statistics Norway (SSB) and add it to the data set. We note that the main subjects of our analysis, Atlantic salmon and Brent Crude oil, are represented by the export price of Norwegian farmed Atlantic salmon (PSALM, USD/kg) and the Europe Brent spot price (BCO, USD/barrel) respectively.

Certain modifications to the data set is needed. We remove indices as to avoid adding collinearity into the model by construction. For the same reason, we remove variables that are highly correlated ($\rho_{i,j} > 0.95$) so that only one representative variable is kept. It is worth mentioning that not all time series begin in 1980, but are included nevertheless, as the estimation procedure allows for time series of uneven length by treating their unreported period as missing values (see Section 3.2). After this clean-up, we are left with an information set consisting of 44 individual commodities. Table (1) lists all the variables along with descriptive statistics and certain statistical tests. The variables are grouped according to their sector affiliation given by the World Bank. Here, Brent Crude belongs to the energy sector, and we add Atlantic salmon to a renamed Animal protein sector.

An important model assumption to accommodate for is that of stationary observable variables. As commodities are known to be I(1) processes, we transform all time series by taking the first difference of log prices to ensure stationarity. All the log differenced time series reject the null hypothesis of the augmented Dickey-Fuller (ADF) test at the 1% significance level, as can be seen in Table (1).

[12] An up-to-date data set can be found at http://www.worldbank.org/en/research/commodity-markets.

Table 1: *Information set $\mathcal{I}^*$ with descriptive statistics.*

Time series	Transf.	Max	Min	Mean	Std	JB[a]	LB[b]	ADF[c]
Animal Proteins								
Atlantic salmon	Δln	0.165	-0.21	0.005	0.069	46844	80.68	-10.31
Beef	Δln	0.13	-0.178	0.004	0.044	49.9	112.99	-9.63
Chicken	Δln	0.12	-0.567	0.005	34.6	98.75	21.12	-8.98
Sheep	Δln	0.079	-0.062	0.005	0.029	20941	57.06	-10.45
Shrimps	Δln	0.212	-0.17	-0.001	0.049	97.28	117.1	-7.41
Beverages								
Cocoa	Δln	0.231	-0.195	0.005	0.062	17380	27.74	-12.38
Coffee. Arabica	Δln	0.269	-0.148	0.003	0.059	48.01	29.2	-11.65
Coffee. Robusta	Δln	0.183	-0.172	0.004	0.056	45474	38.83	-11.16
Tea. avg 3 auctions	Δln	0.141	-0.244	0.001	0.049	87.86	43.59	-12.16
Energy								
Coal, Australian	Δln	0.364	-0.329	0.005	0.073	175.84	66.1	-9.61
Coal, Colombian	Δln	0.234	-0.234	0.004	0.067	154.76	47.62	-6.87
Crude oil, Brent	Δln	0.187	-0.313	0.004	0.09	45.36	37.07	-11.21
Crude oil, WTI	Δln	0.217	-0.324	0.003	0.087	43.9	44.63	-10.55
Natural gas, US	Δln	0.478	-0.393	-0.003	0.138	15.36	18.49	-14.39
LNG, Japan	Δln	0.342	-4.012	0.005	0.125	12.39	27.86	-7.56
Fats & Oils								
Coconut oil	Δln	0.254	-0.26	0.005	0.074	461.43	64.29	-10.7
Copra	Δln	0.266	-0.213	0.006	0.078	262.99	59.34	-10.61
Fish meal	Δln	0.2	-0.11	0.006	0.041	54.07	57.19	-9.24
Groundnut oil	Δln	0.261	-0.21	0.003	0.051	282.39	133.06	-7.39
Groundnuts	Δln	0.14	-0.177	0.005	0.056	114.75	139	-7.11
Palm oil	Δln	0.258	-0.347	0.005	0.069	144.85	50.44	-10.1
Soybean meal	Δln	0.206	-0.186	0.004	0.062	15.26	51.19	-9.97
Soybeans	Δln	0.194	-0.199	0.004	0.057	14.76	48.31	-10.01
Fertilizers								
DAP	Δln	0.365	-0.243	0.002	0.056	110.34	31.22	-9.87
Urea	Δln	0.287	-0.291	0.003	0.078	98.76	37.09	-10.54
Grains								
Barley	Δln	0.233	-0.279	0.002	0.066	53.62	39.4	-9.71
Maize	Δln	0.22	-0.245	0.003	0.061	40.63	41.91	-11.49
Sorghum	Δln	0.254	-0.278	0.003	0.065	64.31	37.11	-12.4
Rice, Thai	Δln	0.35	-0.208	0.005	0.059	447.65	84.91	-8.57
Wheat, US SRW	Δln	0.258	-0.26	0.003	0.071	40.02	33.54	-11.67
Metals								
Aluminum	Δln	0.148	-0.217	0.002	0.05	59.93	65.63	-10.96
Copper	Δln	0.231	-0.35	0.006	0.068	271.95	70.81	-9.01
Lead	Δln	0.24	-0.293	0.008	0.076	66.29	41.85	-11.11
Nickel	Δln	0.248	-0.382	0.003	0.087	432.07	32.49	-10.41
Tin	Δln	0.162	-0.243	0.007	0.064	444.40	47.82	-10.45
Zinc	Δln	0.244	-0.287	0.005	0.067	34.67	46.72	-10.15
Precious Metals								
Gold	Δln	0.112	-0.125	0.008	0.038	27426	20.29	-12.44
Platinum	Δln	0.233	-0.293	0.002	0.058	291.35	53.62	-10.24
Silver	Δln	0.195	-0.214	0.006	0.069	297.68	35.34	-11.4
Raw Materials								
Logs, Malaysia	Δln	0.115	-0.111	0.002	0.031	254.77	81.68	-9.78
Logs, West Africa	Δln	0.144	-0.174	0.003	0.034	214.32	34.96	-11.14
Sawnwood, Malaysia	Δln	0.064	-0.077	0	0.023	24.46	20.99	-11.33
Plywood	Δln	0.068	-0.117	0.001	0.023	170.43	26.84	-11.09
Woodpulp	Δln	0.079	-0.133	0.001	0.035	35.31	140.05	-5.9

[a] Critical values JB: $\chi^2_{2,0.1}$>4.61, $\chi^2_{2,0.05}$>5.99 & $\chi^2_{2,0.01}$>9.21

[b] Critical values LB: $\chi^2_{6,0.1}$>10.64, $\chi^2_{6,0.05}$>12.59 & $\chi^4_{6,0.01}$>16.81

[c] Critical values ADF: $\tau_{0.1}$<-1.62, $\tau_{0.05}$<-1.95 & $\tau_{0.01}$<-2.58

4.3 Methodology

4.3.1 Model

For the purposes of this application, we model each commodity's returns with a *three-factor representation*. Letting $y_{i,t}$ denote commodity i's return at time t, then:

$$y_{i,t} = \underbrace{\lambda_i^g f_t^g}_{\text{global component}} + \underbrace{\lambda_{i,j}^s f_{j,t}^s}_{\text{sectoral component}} + \underbrace{\epsilon_{i,t}}_{\text{Idiosyncratic component}} \tag{8}$$

Each variable is represented by a global factor f_t^g , a sectoral factor $f_{j,t}^s$ (unique for the commodity's sector j), and a commodity-specific factor $\epsilon_{i,t}$ (also referred to as the idiosyncratic factor). The scalars λ_i^g and $\lambda_{i,j}^s$ are factor loadings, and specify commodity i's sensitivity to shocks in the global factor and sectoral factor respectively. These factors capture the co-movement between commodity returns. In the extreme case, a commodity with $\lambda_i^g = \lambda_{i,j}^s = 0$ will have a return that is completely idiosyncratic, displaying no co-variation with other commodities.

In order to obtain the three-factor representation in Equation (8), which will allow us to make interpretations of the factors and the model parameters in a structural sense, we have to impose certain identifying restrictions on the general model introduced in Equations (1)-(3) from Section 3.1, which are restated below:

$$y_t = \Lambda F_t + \epsilon_t \tag{1}$$

$$F_t = A_1 F_{t-1} + ... + A_p F_{t-p} + u_t \qquad u_t \sim i.i.d\ N(0,Q) \tag{2}$$

$$\epsilon_t = \rho\epsilon_{t-1} + e_t \qquad e_t \sim i.i.d\ N(0,R) \tag{3}$$

To accommodate for sectoral blocks, we must impose restrictions on the Λ matrix containing the factor loadings. If variable i does not belong to sector j, we set $\lambda_{i,j}^s = 0$. We rely on the sectors defined by the World Bank when creating the sectoral factors. If we order the variables after their sector affiliation, then Λ becomes a block diagonal matrix.[13] In the general model the factors are modelled as a VAR(p) process, but as there are controversies surrounding the interpretation of VAR models, we model the factors as AR(p) processes for the purposes of this application. This is done to keep the model simple and interpretable. We achieve this by restricting the $A_1, \ldots, A_p$ matrices to be diagonal. We set the order of

[13] More precisely, the part of Λ that contains the sectoral factor loadings will be block diagonal.

the AR processes to one, i.e. $p = 1$. Other non-zero values for p produce similar results. Notice that we have assumed that the factors are uncorrelated, i.e. the factor shocks u_t are i.i.d. We achieve this through restricting the co-variance matrix Q to be diagonal. This implies that we do not allow sector-specific shocks to spill over into other sectors, which ensures that the factors represent pure sector-specific trends. All these restrictions are illustrated in the vector representation of the three-factor representation displayed in Equations (9) and (10).

$$\begin{pmatrix} y_{1,t} \\ y_{2,t} \\ \vdots \\ y_{n,t} \end{pmatrix} = \begin{pmatrix} \Lambda_1^g & \Lambda_1 & 0 & \cdots & 0 \\ \Lambda_2^g & 0 & \Lambda_2 & & \vdots \\ \vdots & \vdots & & \ddots & 0 \\ \Lambda_r^g & 0 & \cdots & 0 & \Lambda_r \end{pmatrix} \begin{pmatrix} f_t^g \\ f_{1,t}^s \\ f_{2,t}^s \\ \vdots \\ f_{r,t}^s \end{pmatrix} + \begin{pmatrix} \epsilon_t^g \\ \epsilon_{1,t} \\ \epsilon_{2,t} \\ \vdots \\ \epsilon_{r,t} \end{pmatrix} \tag{9}$$

$$\begin{aligned} \begin{pmatrix} f_t^g \\ f_{1,t}^s \\ f_{2,t}^s \\ \vdots \\ f_{r,t}^s \end{pmatrix} &= \begin{pmatrix} a_g^1 & 0 & 0 & \cdots & 0 \\ 0 & a_1^1 & 0 & & \vdots \\ 0 & 0 & a_2^1 & & \vdots \\ \vdots & & & \ddots & 0 \\ 0 & \cdots & \cdots & 0 & a_r^1 \end{pmatrix} \begin{pmatrix} f_{t-1}^g \\ f_{1,t-1}^s \\ f_{2,t-1}^s \\ \vdots \\ f_{r,t-1}^s \end{pmatrix} + \cdots \\ &+ \begin{pmatrix} a_g^p & 0 & 0 & \cdots & 0 \\ 0 & a_1^p & 0 & & \vdots \\ 0 & 0 & a_2^p & & \vdots \\ \vdots & & & \ddots & 0 \\ 0 & \cdots & \cdots & 0 & a_r^p \end{pmatrix} \begin{pmatrix} f_{t-p}^g \\ f_{1,t-p}^s \\ f_{2,t-p}^s \\ \vdots \\ f_{r,t-p}^s \end{pmatrix} + \begin{pmatrix} u_t^g \\ u_{1,t} \\ u_{2,t} \\ \vdots \\ u_{r,t} \end{pmatrix} \end{aligned} \tag{10}$$

Here, Λ_j^g and Λ_j are vectors with factor loadings for the variables affiliated with sector j. If the variables i through k, where $1 \leq i \leq k \leq n$, belong to sector j then $\Lambda_j = (\lambda_{i,j}, \lambda_{i+1,j}, \ldots, \lambda_{k,j})'$.

4.3.2 Model Specification

In the three-factor representation presented in Equation (8) we have given the impression that the global component is represented by only one global factor. However, in general, the global component could be represented by any number of global factors (f_t^g). For a given data set, the decision of how many global factors to include needs to take into account the trade-off associated with improved in-sample fit and loss of parsimony that occurs when the number of global factors increases. If we represent the global component by $G > 1$ global factors, we would have that $f_t^g = (f_{1,t}^g, \dots, f_{G,t}^g)'$ and $\lambda_i^g = (\lambda_{i,1}^g, \dots, \lambda_{i,G}^g)$. Several information criteria have been proposed in order to decide the true number of factors to include (Bai and Ng, 2007; Hallin and Liška, 2007; Breitung and Pigorsch, 2012). We use the two information criteria proposed by Bai and Ng (2007), and denote these IC_1 and IC_2. The number of global factors that gives the lowest value of the information criteria are chosen. Table (2) reports the criteria for $\mathcal{I}^*$. Both criteria suggest that one global factor is appropriate.

Table 2: *Information criteria for number of global factors*

	Number of global factors				
IC	1	2	3	4	5
Bai & Ng (1)	-0.053	0.043	0.139	0.235	0.332
Bai & Ng (2)	-0.049	0.051	0.151	0.251	0.352

There is currently no information criterion that gives the optimal number of factors for dynamic factor models with block structures. Therefore, we have not been able to provide any robust procedure for determining the optimal number of factors to represent each sectoral component. We choose to represent each sector component using one factor. As the sectors are subsets of $\mathcal{I}^*$, and thus consist of fewer variables to extract co-movements from, we find it reasonable to represent the sectoral components by the same number of factors or fewer than the global component. We are left with the option of representing the sectoral components by one factor.

4.3.3 Model Estimation

We estimate the model using the EM algorithm as outlined in Section 3.2. The EM algorithm requires us to set the threshold (τ) of the convergence test (Equation (7)), which essentially decides to what degree the model should be fitted to the observed data. A lower threshold results in a more fitted model. The key objective of this application is in fact to provide the best possible fit between the model and the data,[14] as we want to extract the common factors that best describes the co-movement between the commodities. Hence, we set the threshold very low, $\tau = 10^{-7}$.

4.3.4 Variance Decomposition

One way to measure the extent of the common factors' influence on the variables is through variance decomposition. By computing each components' contribution to the total variation in the observed variables, it is possible to state what fraction of a variable's price fluctuations is attributable to a given component (Brooks, 2014). We let ϕ_i^g, ϕ_i^s and ϕ_i^ϵ denote the proportion of the total variation of variable i attributable to the global, sectoral, and idiosyncratic components, respectively. We then have that,

$$\begin{aligned} \phi_i^g &= \left(\lambda_i^g\right)^2 var(f_t^g)/var(y_{i,t}) \\ \phi_i^s &= \left(\lambda_{i,j}^s\right)^2 var(f_{j,t}^s)/var(y_{i,t}) \\ \phi_i^\epsilon &= var(\epsilon_i)/var(y_{i,t}), \end{aligned} \tag{11}$$

where the total variance for variable i is given by:

$$var(y_{i,t}) = \left(\lambda_i^g\right)^2 var(f_t^g) + \left(\lambda_{i,j}^s\right)^2 var(f_{j,t}^s) + var(\epsilon_{i,t}) \tag{12}$$

For the decomposition in Equation (12) to hold, which is a prerequisite for getting ϕ_i^g, ϕ_i^s and ϕ_i^ϵ to sum to one, the factors must be uncorrelated. We ensure that the factors are orthogonal, i.e. uncorrelated, through the restriction of the factor shocks u_t as described earlier.

We remark that we have an approximate dynamic factor model, as mentioned in Section 3. This implies that some of the commodities with large model residuals could potentially have idiosyncratic components that are either weakly correlated

[14] We remark that this is in contrast with the aim of Application B, where we will be concerned with problems of over-fitting.

serially or cross-sectionally. One of the reasons for modelling the idiosyncratic factors as AR(p) processes is to alleviate this issue. As a result, the variance decomposition is considered approximate orthogonal.

4.3.5 Model Interpretation

Previous studies have suggested that the model components provided by factor models can be given interpretations beyond representing different levels of commodity co-movement (e.g., see Alquist and Coibion, 2014; Bilgin and Ellwanger, 2017; Delle Chiaie et al., 2017). It is suggested that the components can in fact be interpreted along a supply/demand dimension, where the global factor mainly captures global commodity demand, while the idiosyncratic factors mainly capture commodity-specific supply shocks. The arguments are as follows.

Due to the way it is constructed, the global factor captures the price trends that are common to all commodities in the data set. The literature discussed in Section 4.1 suggests that demand shocks from global economic activity are likely to affect most commodities by shifting the demand curve in the same manner across markets. It is then argued that this is the movement captured by the global factor. This interpretation has proven to be quite robust, and is for instance strengthened by previous commodity factor model studies reporting that their global factor shows a strong positive correlation with recognized benchmarks for global economic activity such as the Kilian (2009) Index and the Baltic Dry Index (e.g., see Delle Chiaie et al., 2017). However intriguing, this economic interpretation needs to be nuanced somewhat. In their theoretical study, Alquist and Coibion (2014) suggest that the global factor might not solely reflect demand shocks, but potentially also other common influences that feed through to a large share of commodity prices. They illustrate this effect with higher energy prices, that typically lead to higher commodity prices in general as expensive energy increases production costs in most industries. Then again, Baumeister and Kilian (2014) find that the pass-through from shocks to the price of crude oils to other commodities is limited. We will provide our own assessment of this potentially "contaminating" effect. Furthermore, due to increased financialization of commodity markets, common trends might also partly reflect speculative, rather than only fundamental, commodity demand (e.g., see Cheng and Xiong, 2014; Ohashi and Okimoto, 2016).

The sectoral component likely includes both supply and demand forces. Recall that the sectoral factors capture dynamics common to the commodities in the sector they represent. The sectoral components might therefore represent both demand trends common to the sub-group, as well as supply-side shocks spreading within their sector (e.g. through substitution effects). There is no simple way of separating the relative contribution from the two sides.

Opposite to demand shocks, supply shocks tend to be isolated to individual markets or at most to smaller groups of substitutes. Assuming that demand forces are largely captured by the global and the sectoral components, the idiosyncratic component can reasonably be expected to mostly represent supply shocks. However, demand shocks might in some instances be isolated to individual commodities due to factors such as introduction of specific regulations[15], product innovation or technology shifts. The effects of such demand shocks are likely to appear in the idiosyncratic component.

The interpretations (with caveats) presented here, will be made use of when discussing our empirical findings.

[15] E.g. health regulations, tariffs, political sanctions affecting trade.

4.4 Empirical Results

This section is divided into three parts. First, we examine the properties of the global and sectoral factors. Second, we perform variance decomposition analysis on both a broader sectoral level, as well as for Atlantic salmon and Brent Crude oil specifically. Lastly, we perform cumulative return decomposition analysis, and examine historical price shocks for Atlantic salmon and Brent Crude oil in detail, comparing our findings to economic theory and existing literature on the price behaviour of the given commodities.

4.4.1 Estimated Factors and Loadings

Appendix A Figure (14) displays plots for the estimated time series for the global and the 9 sectoral factors, with Appendix A Table (11) displaying the corresponding factor loadings for all commodities in our data set. Combined with the idiosyncratic components, these constitute the input data utilized in the analyses performed in the upcoming sections.

The Global Factor

The global factor time series estimated by our model captures price trends common to all commodities in our data set. In the previous section we introduced the perspective that what this factor mainly reflects, is shocks to global (expected and realized) demand for commodities and the global business cycle. We wish to investigate this proposition empirically.

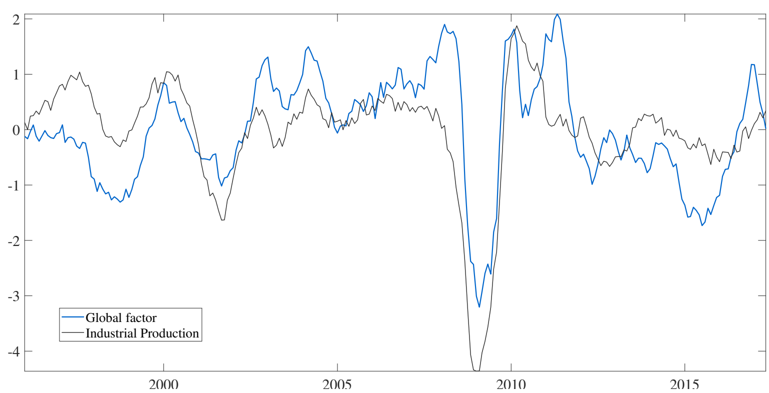

Figure 3: *Global factor v. OECD index of Industrial Production. Both time series are year-on-year and normalized.*

When plotting the global factor, we observe that its peaks and downturns corresponds well with major global economic events, including the large expansion leading up to, and the crash defining, the Global Financial Crisis of 2008. Further, we find a high, positive correlation when testing the global factor against a recognized benchmark for global economic activity. Figure (3) plots the global factor along with the OECD Industrial Production Index. The two time series display a very similar shape and are positively correlated by a coefficient of 0.64. This corroborates the similar findings from Delle Chiaie et al. (2017).

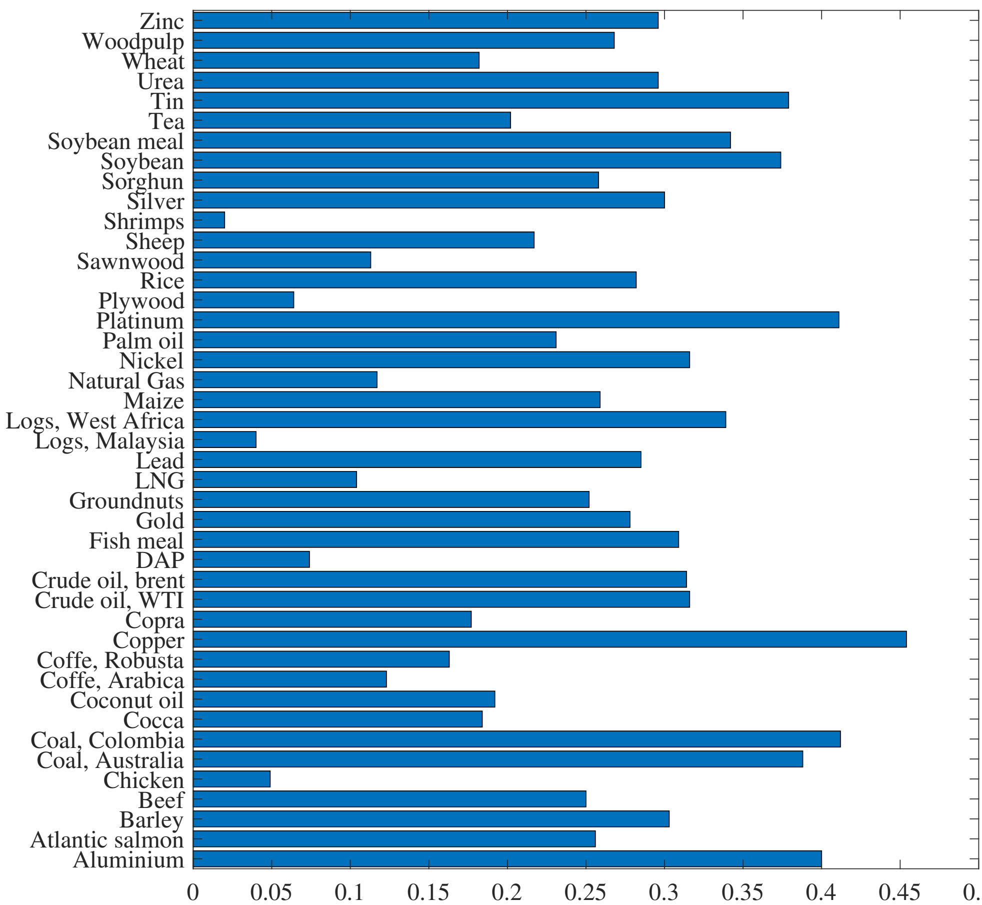

Figure 4: *Global factor loadings. The figure displays each commodity's sensitivity to shocks in the global factor. Note that all commodities display a* positive *relationship with the factor.*

Figure (4) displays each individual commodity's loading on the global factor. Observe that *all* of the World Bank primary commodities have a positive loading on the global factor. Thus; in general, when the global factor increases, the prices in all markets increase. This supports the intuition that an expanding (contracting) global economy increases (decreases) demand for a broad group of commodities - moving prices in the same direction. Further, from Figure (4) we also see that the factor loadings of the energy commodities in the data set are not abnormally high compared to the rest of the commodities in the set. This is an indication that potential energy spill-over effects discussed in Section 4.3.5, are unlikely to be a major driver of the global factor.

In summary; commodity prices consistently increase as the global factor increases, the global factor yields high, positive correlation with several established benchmarks for global economic activity, and the factor does not seem to be driven by other common influences such as energy costs. This makes us confident in mainly interpreting the global factor as representing trends in global commodity demand and the global business cycle.

The Sectoral Factors

Recall that the sectoral factors capture dynamics common to the commodities in its sector, in excess of what is extracted by the global factor. From Appendix A Figure (14), we observe that there are notable differences between the global and the 9 sectoral factors, indicating that they do in fact represent distinctly different movements and play different roles at different points over time.

We note that if a latent factor is extracted from a group of very weakly correlated variables, the likelihood estimation procedure may end up prioritizing the variation in one or a few variables somewhat co-varying with the remaining variables in the group. Highly unequal λ_i^s's within the sector are an indication of this having occurred. While nearly all of our sectoral factors seem to represent genuine sector-wide co-movements, the energy factor is somewhat dominated by the two crude oil benchmarks. This finding will inform our upcoming analysis of Brent Crude.

4.4.2 Variance Decomposition

Estimating variance decompositions is our main tool for assessing the degree of co-movement in the commodity returns. Recall that through this method, we are able to determine what fraction of a variable's price fluctuations is attributable to the global, sectoral and idiosyncratic components. Table (3) displays the sector averages for the variance explained by the global and sectoral components for the full sample period (1980-2018), as well as for two sub-sample periods (1980-1999 and 2000-2018).

Table 3: *Variance decomposition of returns (ϕ_i^g and ϕ_i^s), sector averages.*

	Global component			*Sectoral component*		
	Full sample	1980-1999	2000-2018	Full sample	1980-1999	2000-2018
All	8.73%	4.29%	17.10%	27.97%	30.31%	24.81%
Animal proteins	5.57%	5.92%	9.38%	11.32%	10.14%	5.16%
Beverages	3.57%	0.56%	8.73%	32.65%	37.86%	23.75%
Energy	11.32%	4.48%	21.48%	42.04%	49.98%	34.09%
Fats & oils	7.89%	9.34%	15.70%	28.00%	30.77%	24.72%
Fertilizers	6.22%	2.85%	4.22%	34.23%	9.61%	46.72%
Grains	8.27%	3.62%	12.68%	29.87%	33.23%	28.60%
Metals	15.87%	2.13%	40.01%	19.58%	21.43%	20.98%
Precious metals	13.70%	3.43%	24.53%	43.71%	56.96%	40.28%
Raw materials	5.02%	3.41%	7.95%	26.33%	28.31%	22.83%

Table (3) reports that global shocks account for a moderate fraction of commodity price variability. Over the full sample period, an average of 8.7% of the price fluctuations in the commodities in our data set is attributable to the global factor. For the sectoral factors, this number is 28%. Thus, global and sectoral shocks account for about a third (36.7%) of fluctuations in commodity prices, leaving two-thirds to being explained by commodity-specific forces. We notice that the commodity groups most responsive to global shocks, are the two metals sectors.

We also estimate our model separately for the time periods of 1980-1999 and 2000-2018. Interestingly, the sub-sample analysis reveals that there has been a significant change in the relative importance of the three components. Between 1980-1999, the average variance explained by the global, sectoral and the idiosyncratic components is 4.3%, 30.3% and 65.4% respectively. For the second sub-sample the equivalent numbers are 17.1%, 24.8% and 58,1%. Notably, the global factor's explanatory power has increased significantly across commodities and sectors over the last decades. Not only does the total average variance explained by the global components greatly increase (up 12.8%) between the sub-samples, we do in fact observe that the global components' importance has increased significantly for every single sector. We interpret this as a clear indication of more integrated commodity markets. In other words, our findings support the consensus from the literature discussed in Section 4.1.[16]

Atlantic salmon and Brent Crude oil

Moving from general to commodity-specific insights, we now use variance decomposition to examine the price fluctuations of Atlantic salmon and Brent Crude oil specifically.

Table 4: *Variance decomposition of commodity returns for Atlantic salmon and Brent Crude oil (ϕ_i^g and ϕ_i^s).*

	Global component			*Sectoral component*		
	Full sample	1980-1999	2000-2018	Full sample	1980-1999	2000-2018
Atlantic salmon	7.99%	4.98%	9.17%	1.37%	2.48%	1.90%
Brent Crude oil	12.45%	1.97%	28.54%	73.88%	86.98%	60.94%

Table (4) reports that shocks to the global factor contributes moderately to the variation in the Atlantic salmon time series. Specifically, the variance decomposition shows that Atlantic salmon's global component explains 7.99% of the price fluctuations over the sample period, which can be interpreted as 8% of the variability in Atlantic salmon returns over this period being due to effects from global commodity trends.[17] Interestingly, the sub-samples indicate that this relationship has changed significantly over the last 40 years, with variance explained increasing from 5.0% to 9.2% between the sample windows. We conclude that although the relationship is moderate, Atlantic salmon can not be said to be disconnected from

[16] Most notably Cheng and Xiong (2014) and Juvenal and Petrella (2015). These papers mainly attribute this development to demand shocks from the unexpected high growth in emerging markets, along with the emergence of commodity index investment starting around 2004.

[17] As an ancillary metric, we can examine correlation. The correlation between the Atlantic salmon time series and the global factor is *0.26* over the full sample period, indicating that the Atlantic salmon returns tend to move in the same direction as the global factor.

broader trends in global commodity demand or the real economic cycle.[18]

Regarding the sectoral component, it has proven difficult identifying any significant co-movement between Atlantic salmon and other animal proteins, in excess of that captured by the global factor. It would have been interesting to assess the co-movement between Atlantic salmon and a pure seafood sector block. However, shrimp is the only seafood commodity in the World Bank database.[19] In any case, we note that the model does detect some common movement within the Animal protein sector, but that Atlantic salmon's sector loading is fairly low. We also note that any significantly lagged relationships between the commodities, including some substitution effects, will not be captured in this modelling of co-movement of returns. There are however other relationships one could have hypothesized to be present among the set of animal protein commodities. An example could be events of larger protein consumption elasticity in response to changes in real economic activity, compared to the full set's average commodity demand elasticity.

Atlantic salmon stands out as a commodity with highly idiosyncratic price behaviour, with the variance decomposition indicating that it is shocks mostly isolated to the Atlantic salmon market that is driving the vast majority of the price movement. The large idiosyncratic component represents shocks related to commodity-specific supply, as well a the potential existence of idiosyncratic demand. The salmon supply's inherent dependence on biological, environmental and seasonal factors, makes it a natural contributor to the commodity's idiosyncratic price fluctuations.

Moving on to Brent Crude, we find the following relationships. The variance decomposition shows that its global component explains 12.45% of the variability over the sample period. Further, the vast majority of the remaining variation in the Brent price is attributable to its sectoral component with 73.88% of variance explained. This isn't very surprising, given the fairly tight level of integration between the markets for Brent and WTI benchmark oils, combined with the observation that these commodities somewhat dominate our model's energy sectoral factor.[20] Most interesting, however, are the findings from the sub-sample analysis

[18] We note that market financialization resulting from the 2006 introduction of the Fish Pool ASA futures market, might have contributed to Atlantic salmon's increased co-movement with global commodity markets.

[19] A solution would of course be to introduce external time series to our data set. Unfortunately, we have been unable to acquire suitable time series for relevant seafood commodities spanning the long sample period.

[20] Observing the energy factor in Appendix A Figure (14), we clearly see that major jumps in the factor corresponds to major historical events assumed largely isolated to the oil markets.

regarding the global component. Namely, we observe that the variance explained by the global component has increased from 1.97% in 1980-1999 to 28.54% in 2000-2018. So while the vast majority of the variation in Brent Crude prices up until 2000 was attributable to sectoral and commodity-specific shocks, since 2000 close to 30% of the fluctuation is explained by the commodity's sensitivity to shocks to the global factor. Corroborating the findings from the variance decomposition, the correlation between the time series has also increased significantly the last decades, from 0.27 between 1980-1999 to 0.53 between 2000-2018.

In summary we find that although mostly defined by sectoral or commodity-specific shocks, the price fluctuations of both Atlantic salmon and Brent Crude are to a significant extent formed by common trends in global commodity markets. Furthermore, for both commodities, the global factor has increased substantially in importance over the last decades.

4.4.3 Cumulative Return Decomposition

The properties of logarithms allow us to simply add monthly log returns in order to find the total return over a longer time period. As the global, sectoral, and idiosyncratic components in the three-factor representation (Equation 8) are log return time series, we can accumulate their returns to obtain the total return of each component for a given time period. These components' accumulated returns will together add up to the total return of the given commodity. With this decomposition we are able to assert what proportion of the total return over a specific period can be attributed to global, sectoral, and idiosyncratic effects.

This allows us to zoom in at specific price events and answer questions such as: which factors contributed the most to a given price movement? Are there any forces pulling the price in the same or opposite direction? The upcoming subsections are dedicated to cumulative return decomposition of specific, noteworthy price events for Atlantic salmon and Brent Crude.

Before we delve into the price event analysis, we will give a brief introduction and explanation of the cumulative return plot. Figure (5) shows the cumulative return of Atlantic salmon and Brent Crude from 2004 through 2018. The black solid lines represent the total return accumulated since 2004. Correspondingly, the coloured bars indicate the total accumulated returns attributable to the global, sectoral and idiosyncratic components up until each point in time. The coloured bars are stacked, so that they together add up to the total return.

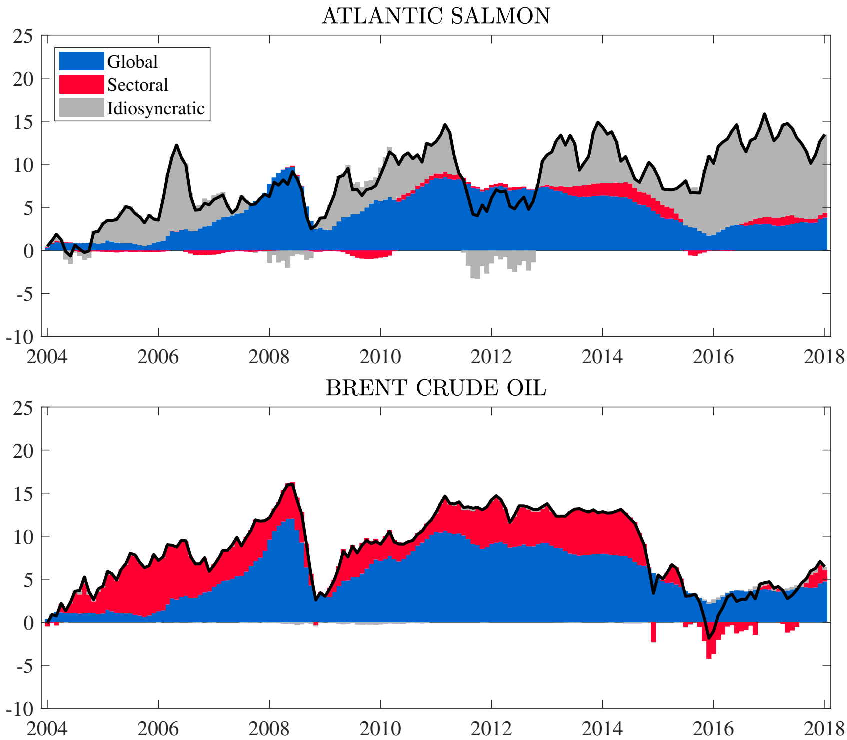

Figure 5: *Return decomposition from 2004:M1 - 2018:M3. The figure presents decompositions of cumulative normalized log returns for Atlantic salmon and Brent Crude oil.*

Several aspects of the three-factor representation can be observed through Figure (5). If we examine the global components (in blue), we observe that they are identical in shape; the global components are just scaled versions of the global factor (f^g), as each variable has a unique loading (λ_i^g's) on the global factor. Further, as Atlantic salmon and Brent Crude belong to different sectors, their sectoral components (in red) have different shapes. It is worth noting that relative to the total accumulated return (black solid line) they also vary in scale, indicating the relative importance or contribution of the sectoral factor to the return of the respective commodities. Lastly, we note that the idiosyncratic components (in grey) capture the remaining, unexplained part of the returns.

Next we will examine noteworthy price events for Atlantic salmon and Brent Crude, and gauge our model output against the existing literature.

Event analysis, Atlantic Salmon

Figure (6) displays cumulative return decompositions of four historical episodes of particular interest with respect to the Atlantic salmon market.

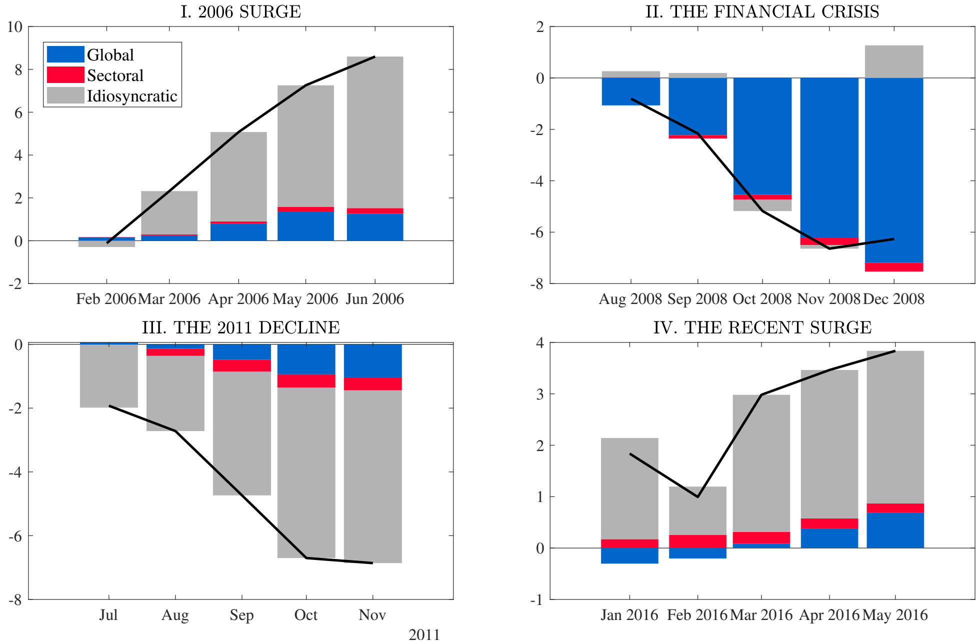

Figure 6: *Return decomposition of Atlantic salmon price events (cumulative normalized log returns).*

We observe an incredible surge in the salmon price between primo and medio 2006. Our decomposition illustrated in Figure (6.I) indicates that this shock was driven by mostly idiosyncratic forces. Looking to the literature, supply-limiting factors such as adverse weather conditions and disease problems in Chile in early 2006, can help explain the tight market and high prices displayed that year, according to Globefish (2006). Oglend (2013) further illuminates potential supply-side causes for this price movement. He points to the fact that if biomass is lower than expected prior to the most potent production period, farmers must be compensated by higher prices in order to give up the valuable high growth period. And that in such circumstances there is not enough fish to satisfy both consumption and "production" demand. Players who want fish on the spot market must therefore bid up the price. These spikes appear on occasion, and did so in the spring of 2006, Oglend (2013) argues. What about potential demand-side factors? Actually, Brækkan and Thyholdt (2014) find the second largest positive demand shift (20%) in their sample period (2002-2011) from 2005 to 2006. Comparably, we find that the global component and sectoral compents contributed moderately (15%) to the observed price spike in this time period.

Excluding the 2006 spike, from 2004 to medio 2008 we observe a trend of a gradual increase in the salmon price that seems to be largely driven by a steady increase in the global factor (see Figure (5)). This increase matches the world macroeconomic expansion that started in the early 2000s and lead up to the 2008 financial crisis. And, as illustrated in Figure (6.II), in the second half of 2008 we see that a dramatic fall of the global factor, representing an across-the-board rapid decline in commodity prices, causes the salmon price to fall. In this time interval, we are comfortable in asserting that the movement in the Atlantic salmon price was almost exclusively driven by shocks to global economic activity.

In 2011 we see a large decline in the salmon price, and it remained low until 2013. The return decomposition illustrated Figure (6.III) attributes most of the decline to idiosyncratic forces, though there is also a significant contribution from the global and sectoral components (22%). One event worth discussing here is the Chilean salmon crisis. Starting in late 2007 the Chilean salmon industry experienced severe disease issues. This had a significant negative effect on production from 2009 until at least early 2011, leading to an increased supply/demand gap. However, from 2011 the supply increased significantly. Even taking in to consideration the somewhat limited overlap with Norwegian export markets, this increased supply from Chile undoubtedly contributed to the observed price shock in 2011, and the sustained lower prices the following years. An idiosyncratic supply shock, that is. Simultaneously, Brækkan and Thyholdt (2014) finds a significant demand drop in the EU for Norwegian salmon in 2011. This demand shock corresponds well with an observed decline in the global factor in the same time period, indicating that this drop was part of a general trend.

Starting in late 2015, another surge in the spot price for salmon occurred. Most industry experts explain the recent surge based on increased global demand for salmon, combined with a limited ability to increase supply due to stringent regulations, and parasite and disease issues. These supply-limitations will naturally be isolated to the salmon industry, and are therefore in line with the very large idiosyncratic component estimated in this time period, displayed in Figure (6.IV). An example of such a factor is the toxic algae bloom appearing in Chile in this period, leading to Chilean production volumes falling by almost 20% (Terazono, 2016). Regarding the argued increase in demand for salmon, if this was in fact a major price driver, we can categorize it as being of a largely idiosyncratic nature, given that we don't observe any significant increase in the general demand for commodities.

Event analysis, Brent Crude Oil

Figure (7) displays cumulative return decompositions of four episodes of particularly large shocks to the price of Brent Crude oil.

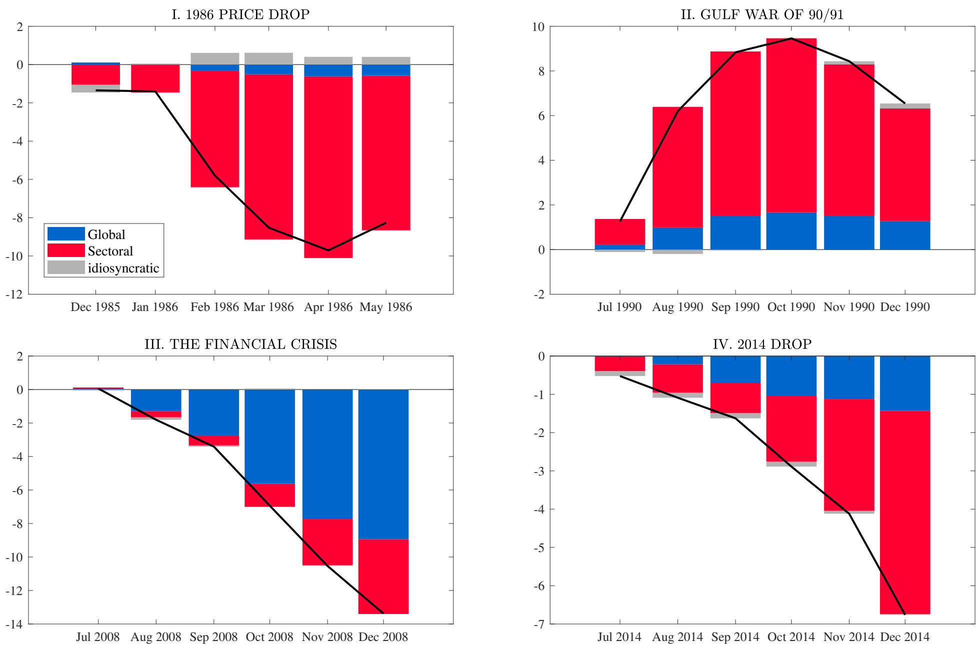

Figure 7: *Return decomposition of Brent Crude oil price events (cumulative normalized log returns).*

Figure (7.I) and Figure (7.II) display the two largest oil price shocks that occurred in the late 1900s. Specifically, the 56% cumulative decline in the price of oil in early 1986, occurring after Saudi Arabia abandoned its policy of stabilizing the price of crude oil, and the spike occurring at the outbreak of the Gulf War in the early 1990s. The return decompositions suggest that sector-specific factors were in fact essentially the sole drivers of these price events.

Referring back to the Brent Crude graph in Figure (5), we see that it suggests that a surge in the global factor was the primary driver for the strong rise in oil prices in the run-up to the Global Financial Crisis. This is in line with the previously addressed narrative that the rise of emerging-market economies in Asia lead to a series of positive demand shocks in the global oil market (Kilian, 2009). Correspondingly, Figure (7.III) shows that the 2008 price drop associated with the GFC was mainly driven by a drop in the global factor. Note that the observed contribution from the sectoral component could indicate that the energy/oil sector was more strongly impacted by this crisis than global commodity markets in general.

Figure (7.IV) displays the drop in the oil price that occurred between June and December 2014, hitting the Norwegian economy hard. We find that around 20% of the negative return in this period is attributed to the global component, while the rest is attributed to the sectoral component. This closely matches the findings of Baumeister and Kilian (2016), who attribute $11 of the $49 decline to the cumulative effects of negative demand shocks, tracing this to a slowing global economy. They find the remaining decline being due to positive shocks to current oil production, as well as an unpredictable shock to oil price expectations in July 2014 which lowered the demand for oil inventories.

In summary, the findings from the event analyses performed for Atlantic salmon and Brent Crude oil above, demonstrate that our DFM can assist in the discussion of sources of major price events for the two commodities.

4.5 Concluding Remarks

In this study, we have investigated the co-movement in global commodity markets and decomposed individual commodity price movements into global, sectoral, and idiosyncratic components. By considering a large set of globally traded commodities retrieved from the World Bank, we obtained a global factor by isolating the movement common to all commodities. By grouping related commodities together, and then isolating the co-movement within each group, we obtained the sectoral factors. The remaining commodity-specific movement was captured by idiosyncratic components. In our analyses, a particular focus has been put on the major Norwegian export commodities of Atlantic salmon and Brent Crude oil. The study yielded some interesting empirical results.

We find empirical support for an interpretation of the global factor as mainly reflecting trends in global commodity demand. Through variance decomposition analysis, we find that a moderate, but significant, share of the fluctuations in the price of Atlantic salmon and Brent Crude oil over the past 40 years can be ascribed to shocks to this factor. Most notably however, a sub-sample analysis reveals that the global factor's importance has significantly increased for both commodities, with ϕ^g_{salmon}=9% and $\phi^g_{bco} = 29\%$ since 2000. Interestingly, the same trend can in fact be seen across all commodity sectors. This indicates that Atlantic salmon and Brent Crude have taken part in a general trend of increased integration between global commodity markets. Regarding sectoral and commodity-specific movement, we find that while monthly Atlantic salmon returns are of a highly idiosyncratic nature, Brent Crude displays stronger co-movement with its sector. We do however note that the energy sector factor estimated could be considered as largely representing trends in the markets for oil more so than truly sector-wide energy commodity trends.

Further, we find that our model forms a crude, but effective, tool for interpreting specific commodity price events. The model output can be used to scrutinize existing narratives on commodity price shocks, and identify relationships unobservable from analysis of individual time series. We believe a DFM approach to analysis of commodity price composition can hold several advantages over other commodity-specific structural models. Conventional structural models typically rely on consumption and production data, but these are often not trivial to acquire, are released with significant time lags or do not even exist for many commodities. This DFM methodology, on the other hand, only relies on publicly available commodity price data.

Major benefits from understanding the composition of commodity price movements relate to macroeconomic concepts influencing fiscal and monetary policy stances, as well as industry and finance considerations regarding risk, hedging and investment decisions. We argue that a DFM approach to decomposing commodity prices should be considered a valuable addition to the existing toolkit used for examining these markets.

5 Application B: Forecasting

Application A demonstrated the DFM framework's ability to decompose price movements and provide insights into the broader sources of these movements. In Application B we switch our focus to the future - to the issue of forecasting commodity spot prices. We use the dynamic factor model framework to produce price forecasts for Atlantic salmon and Brent Crude oil at multiple horizons. This is achieved through constructing a comprehensive set of predictors for both commodities, carrying out a model selection procedure assisted by a genetic algorithm, and finally estimating the dynamic factor models.

Insights into the future price trajectory of a commodity is of obvious interest to a wide range of players. Participants along all parts of the value chain make important decisions regarding their strategy, operations and risk management on the basis of what they believe the future price paths of relevant commodities will be. With regards to our selected commodities, more accurate forecasts of their highly volatile prices would be valuable to stakeholders ranging from E&P companies and salmon farmers, to refineries and salmon processors, not forgetting end-consumers, regulators and financial players.

Section 5 is organized as follows. First, we review literature relevant to the forecasting of Atlantic salmon and Brent Crude oil. Second, we present the information sets used to forecast the spot prices of the two commodities, before we explain our methodology approach in detail. Finally, we assess the forecasting performance of the DFMs against other benchmark models and look at how the DFMs measure up against results found in the literature.

5.1 Literature Review

Insights from the literature on the characteristics and dynamics of the two commodities were provided in Section 2. In Section 3, the essential literature on the general DFM framework was addressed. This review therefore limits itself to discussing literature directly related to the forecasting of the spot prices of Atlantic salmon and Brent Crude oil.

Atlantic salmon

Accurate spot price forecasts are of obvious interest to salmon market participants. However, a search for studies that directly forecast the salmon price produces surprisingly few results. To our knowledge, only three other studies have been published in the last 20 years. Most recently Sandaker et al. (2017) predicted the conditional distribution of weekly prices from 2006 to 2017 using a quantile regression model. A large part of their study is dedicated to finding key predictors of the spot price through the use of a genetic algorithm. However, it is unclear how well their models perform at point-forecasting the median price out-of-sample, which makes their results difficult to compare against. Further, Bloznelis (2018) compares 16 different forecasting models in his comprehensive study of short-term price forecasting. Interestingly, none of the models significantly improves upon the naïve random walk model. This could be due to a very limited set of predictors. Lastly, Guttormsen (1999) proposes six potential forecasting techniques in his quest to find a simple and reliable method to assist salmon farmers in their decision-making. With a Classical Additive Decomposition (CAD) he produces results of correctly predicted direction between 70% and 90% of the out-of-sample cases - however providing no evidence for a general superior model. Moreover, with the salmon market having evolved since then, his study from 1999 could be considered somewhat outdated.

Brent Crude oil

In contrast to Atlantic salmon, direct oil price forecasting is widely explored in the literature, with a wide range of methods utilized. A few studies, arguably the most relevant, are addressed here. Most recently, Westgaard et al. (2017) identified key explanatory variables with a forward selection algorithm and forecasted monthly oil returns with a predictive regression model. Their results suggest that financial variables are superior in predicting the Brent spot price. It is however unclear whether their predictive regression outperforms the random walk on a general basis. Beckers (2015) explores a reduced form VAR model with seasonal "dummy" variables by including explanatory variables such as the Kilian Index, interest spreads, oil production and inventories. He finds that the VAR outperforms the random walk for longer horizons, but is not generally better for shorter horizons. However, a few studies have managed to outperform the no-change benchmark for monthly forecasts; Baumeister and Kilian (2012) report mean squared prediction errors (MSPE) that improve 18% and 25% on a no-change model at 1-month and 3-month forecast horizons, respectively. They achieve this using a VAR model on a combination of futures and commodities. Chen (2014) uses monthly data to investigate the predictive content of oil-sensitive stock price indices for spot crude oil prices using a predictive regression model. Using the NYSE Arca Oil Index as a predictor, he is able to improve 24% on the no-change model's MSPE at the 1-month horizon. Common to these studies is that their models perform well at shorter horizons (1-3 months), both in absolute terms and relative to benchmark models.

This study in context of the existing literature

To our knowledge, ours is the first study utilizing dynamic factor models to forecast price changes for either Atlantic salmon or Brent Crude oil. Therefore, we certainly provide new knowledge on the abilities of such state-space approaches to forecast commodity price movements. The studies discussed above use more or less similar predictor variables in their forecasting models - variables we will adopt in our analysis. However, in contrast to the other studies, we will utilize the variables indirectly. That is, our predictors will instead be latent, unobservable factors extracted from, and based on, the information and co-movement found within sets of variables recognized as key predictors. With our study, we test whether modelling based on such latent factors can improve upon the price change forecasts from the existing literature.

5.2 Data

In our forecasting, we use monthly data in the time period from January 2006 to April 2018. We create separate information sets for each of the two commodities. Each set includes demand-side, supply-side, and other/financial predictors along with the variable to be forecasted. The predictors we include for Atlantic salmon are largely based on the work of Sandaker et al. (2017), who identify superior predictors through a comprehensive variable selection process. The predictors we include for Brent Crude are largely based on the variables identified by Westgaard et al. (2017) and Olimb and Odegaard (2010). These studies follow a systematic and thorough approach as they identify and provide justifications for superior predictors. We supplement with the NYSE Arca Oil Index identified by Chen (2014), found to have substantial predictive value.

Table 5: *List of predictor variables, Atlantic Salmon.*

ID	Name	Trans.	Lag(s)	Window	Effect
CEU	Consumption, Europe	Δln	3, 6, 12	YoY	⇑
CEM	Consumption, Emerging markets	Δln	3, 6, 12	YoY	⇑
EURNOK	Currency pair, EUR/NOK	Δln	3, 6, 12	MoM	⇑
CLDUSD	Currency pair, CLD/USD	Δln	3, 6, 12	MoM	–
TRO	Norwegian trout, spot price	Δln	1-6	MoM	⇑
MAC	Norwegian mackerel, spot price	Δln	1-6	MoM	⇑
SHR	Shrimp (Mexican), spot price	Δln	1-6	MoM	⇑
MEA	Meat price index	Δln	1-6	MoM	⇑
BEF	Beef (U.S.), spot price	Δln	1-6	MoM	⇑
CHI	Chicken (U.S.), spot price	Δln	1-6	MoM	⇑
FED	Feed consumption (NOR)	Δln	1-6	YoY	⇓
LIC	Sea lice occurence (NOR)	Δln	3, 6, 9, 12, 15	YoY	⇕
SEA	Average sea temperature (NOR)	Δln	3, 6, 9, 12, 15	YoY	⇓
SMO	Smolt release (NOR)	Δln	3, 6, 9, 12, 15	YoY	⇓
BIO	Standing biomass, salmon (NOR)	Δln	3, 6, 9, 12, 15	YoY	⇓
HVS	Harvest volumes, salmon (NOR)	Δln	3, 6, 9, 12, 15	YoY	⇓
HVT	Harvest volumes, trout (NOR)	Δln	3, 6, 9, 12, 15	YoY	⇓
FSM	Fish meal (Any origin), spot price	Δln	1-6	MoM	⇑
RAP	Rapeseed oil (Malaysia), spot price	Δln	1-6	MoM	⇑
SOY	Soybeans (U.S.), spot price	Δln	1-6	MoM	⇑
WHE	Wheat (U.S.), spot price	Δln	1-6	MoM	⇑
PSALM	Norwegian farmed Atlantic salmon, spot price	Δln	1-6	MoM	⇑
OSLSFX	OSLO Seafood Index	Δln	1-6	MoM	⇑
MHG	Marine Harvest Group stock	Δln	1-6	MoM	⇑
BCO	European Brent spot	Δln	1-6	MoM	⇕

In order to satisfy the assumption of stationarity, most of the variables need to be transformed. All prices and indexes are transformed by taking the difference of log levels. All rates and spreads are kept at level. The remaining variables are kept at level if they are stationary, and are log-differenced otherwise. We identify seasonal variables by examining each variable's autocorrelation function. There are several variables that exhibit yearly seasonality, most notably *average sea temperature* and *sea lice occurence*. For these variables we use a year-over-year (YoY) time window; for the rest we use a month-over-month (MoM) time window.

We add auxiliary variables in the form of lagged variables to each information set. Where literature have studied lagged relationships between predictors and the variables of interest, we include the lags recommended. For all other predictors we include variables lagged 1 to 6 months. For Atlantic salmon, the final information set contains 25 unique predictors and 136 variables in total. For Brent Crude oil, the final information set contains 25 unique predictors and 150 variables in total.

Table 6: *List of predictor variables, Brent Crude oil.*

ID	Name	Trans.	Lag(s)	Window	Effect
MSP	US money supply	Δln	1-6	MoM	⇑
UEM	US unemployment rate	-	1-6	-	⇑
GDP	World GDP index	Δln	1-6	MoM	⇑
USD	Trade weighted USD	Δln	1-6	MoM	⇕
TBIL	US 3 month treasury bill rate	Δln	1-6	MoM	⇓
AGR	Agricultural commodity index	Δln	1-6	MoM	⇑
N_GAS	U.S. Natural gas spot price	Δln	1-6	MoM	⇑
IMP	US net petroleum import volumes	Δln	1-6	MoM	⇓
KIL	Kilian Index	-	1-6	-	⇑
OPEC_PR	OPEC production	Δln	1-6	MoM	⇓
OPEC_SC	OPEC surplus capacity	Δln	1-6	MoM	⇓
NOPEC_PR	Non-OPEC production	Δln	1-6	MoM	⇓
RIG	Rotary rigs in operation	Δln	1-6	MoM	⇓
OECD_INV	OECD inventories	Δln	1-6	MoM	⇓
OPEC_INV	OPEC inventories	Δln	1-6	MoM	⇓
US_INV	U.S inventories	Δln	1-6	MoM	⇓
REF	US refinery utilization	Δln	1-6	MoM	⇕
BCO	Europe Brent spot price	Δln	1-6	MoM	⇑
OSX	PHLX Oil Service Index (OSX)	Δln	1-6	MoM	⇑
NYME	NYSE Arca Oil Index	Δln	1-6	MoM	⇑
F_SP	NYMEX futures spread ratio	-	1-6	-	⇕
BB_SP	Yield of BoA Merrill (BB) Index	-	1-6	-	⇓
CRA	NYMEX crack spread	-	1-6	-	⇑
MET	Metal index	Δln	1-6	MoM	⇑
GLD	Gold spot price	Δln	1-6	MoM	⇑

The full lists of predictors used for Atlantic salmon and Brent Crude can be found in Table (5) and (6), respectively. In the table, we also list the transformation applied to each predictor, and where we difference the variable we list the time window (MoM/YoY) used. Additionally, we specify which lags of each predictor is included. Finally, we provide the hypothesized direction of change in the forecasting subject corresponding to a positive change in the given predictor variable. For descriptive statistics and statistical test, see Table (12) and Table (13), both found in Appendix B.

5.3 Methodology

5.3.1 Design of Forecasting Exercise

Our forecasting strategy for this application is as follows. We start with an initial in-sample period spanning January 2006 to March 2015, a total of 112 months. As the time series extend to April 2018, this leaves us with an out-of-sample period of 36 months. This decision is simply a trade-off between model calibration and performance evaluation.[21] In order to realistically replicate the situation facing the forecaster, we follow an *expanding window* approach. That is, after obtaining forecasts for a given horizon we expand the in-sample data to include new information before producing the next forecast. This procedure is repeated until the hold-out period is exhausted. Finally, we produce the actual forecasts using an *"iterated" h-step-ahead* forecasting strategy. That is, we make use of the one month-ahead forecast to produce the two month-ahead forecast, and so on. The alternative would be to employ a horizon-specific approach, i.e. specifying and estimating separate models for each of the forecast horizons. The advantages of our approach are simplicity and seemingly better forecast accuracy if the model is correctly specified (Marcellino et al., 2006). However, creating one model that performs superiorly for a wide range of horizons is close to impossible. We limit our study to 1-3 month-ahead forecasts.

5.3.2 Model

We use the general DFM introduced in Section 3.1 as a basis for the model presented here. For the purpose of forecasting we model the variables we seek to forecast as a linear combination of factors $f_{j,t}$, where $j \in [1, \dots, r]$. Letting $y_{i,t}$ denote the return of variable i at time t, we have that:

$$y_{i,t} = \lambda_{i,1} f_{1,t} + \lambda_{i,2} f_{2,t} + \cdots + \lambda_{i,r} f_{r,t} + \epsilon_{i,t} \tag{13}$$

[21] A reasonable share of the observations should be dedicated to model calibration as we want our model to increase its ability to separate noise from genuine dynamics in the data. At the same time, too small an out-of-sample period leads to difficulties associated with statistical inference of forecasting results.

The unobservable factors $f_{j,t}$ represent common variation found in groups of predictors. By modelling the factors as a VAR(p) process we are able to forecast the factor values. With forecasted factor values we can provide forecasts for all the observed variables. If we let $\hat{f}_{j,t+h}$ denote the forecasted value of factor j, then we find the h-month-ahead value of variable i, denoted $\hat{y}_{i,t+h}$, by finding the expected value of Equation (13):

$$\hat{y}_{i,t+h} = \lambda_{i,1}\hat{f}_{1,t+h} + \lambda_{i,2}\hat{f}_{2,t+h} + \cdots + \lambda_{i,r}\hat{f}_{r,t+h} \tag{14}$$

We note that being able to forecast the predictors themselves is merely a by-product of utilizing a dynamic factor model. The information set, model set-up and methodology are all tailored to the variables of interest (Atlantic salmon and Brent Crude), as it is the model's ability to forecast these variables that we wish to optimize and evaluate.

The factor representation described above is fairly general, and we make few modifications to the general dynamic model framework introduced in Equations (1)-(3), that are restated below:

$$y_t = \Lambda F_t + \epsilon_t \tag{1}$$

$$F_t = A_1 F_{t-1} + \cdots + A_p F_{t-p} + u_t \qquad u_t \sim i.i.d\ N(0, Q) \tag{2}$$

$$\epsilon_t = \rho\epsilon_{t-1} + e_t \qquad e_t \sim i.i.d\ N(0, R) \tag{3}$$

The factors in this model are all so-called block factors, i.e. factors that are common only to subsets of the variables. It would of course be possible to specify global factors here as well (factors representing co-movements common to *all* variables). However, as we are dealing with a quite diverse set of variables we do not expect to find global factors particularly useful. Thus, we only use block factors. To achieve this we have to impose restrictions on the Λ matrix that allow us to specify which variables should load on a given factor; that is, if variable i should not load on factor j, then we restrict $\lambda_{i,j} = 0$. We note that specifying which variables should load on a given factor is equivalent to specifying which variables a factor should extract co-variation from. The block factors can be created so that they only capture the common variation among certain variables, e.g. all variables related to supply-side effects. The factors are modelled as a VAR(p) process in which the factor shocks u_t are free to co-vary. We set the order of the VAR process to six, i.e. $p = 6$.

Modelling the dynamics with such a high degree of freedom enables us to provide forecasts over longer horizons.

For a matrix representation of the Equations (1) and (2) for the model used in Application B, we refer to the matrix representation found in Equations (4) and (5) in Section 3.

We note that there are no information criteria available to specify the number of factors to include in the case of a restricted loading matrix. As we are not able to make a more informed decision, we find it sufficient to represent the co-movement in any subset of variables by one factor only.

5.3.3 Model Selection

While DFMs were primarily developed as an approach to efficiently making use of the rich information available within large-scale data sets, the extraction of dynamic latent factors can be considered a valuable trait regardless of data set size. In fact, for our forecasting purpose, quick empirical testing of different model set-ups has indicated that we do not seem to obtain improved forecasting performance by adding on large numbers of predictors. With this in mind, we have instead chosen to employ a more sophisticated model selection procedure, where we focus more on the benefits of latent factors than on the inclusion of high numbers of predictor variables.

Optimal model selection refers to the task of selecting the superior model among many candidate models. In our case, the number of candidate models is large and equal to all possible combinations of block structures given the data set of potential variables. Quick computations show that the search space of candidate models is vast, and it would be computationally infeasible to do an exhaustive search.[22] We take two measures in order to reduce the size of this optimization problem. First, we restrict the search space to a limited set of latent factors. Second, we employ a genetic algorithm (GA) to select an optimal model from the restricted search space.

We restrict the search space by specifying the maximum number of factors that may be included in the model and also by restricting the maximum number of variables in each block. We set the maximum number of factors to six, and allow the factors to extract co-variation from a maximum of six variables. We let the dependent variable load on all factors as this encourages the factors to find co-

[22]With 100 variables and 10 factors, the loading matrix would contain 1000 elements. Restricting the variable is a binary decision; hence, there are $2^{1000} \approx 10^{30}$ possible combinations

movements between the given variable and its predictors. Setting these restrictions essentially reduces the problem to one of variable selection.

Variable selection, or variable subset selection, is the process of selecting a subset of variables to be used in model construction. For each factor, we wish to select six variables from the information set. We include the option of selecting an "empty" variable, which essentially gives the option of selecting fewer than six variables. This selection procedure is needed for each factor. The search space is considerably reduced, but is still large, so an exhaustive search is still computationally infeasible.

We choose to employ a GA to select an optimal model from the remaining search space. Here, we limit ourselves to shortly explain the principles behind the GA and how we make use of it in our application. For a thorough introduction to GAs and their applications we refer to Mitchell (1996).

A GA is a heuristic that is used to find or generate a sufficiently good solution to an optimization problem. It is inspired by the process of natural selection or "survival of the fittest", and borrows terminology from biology. Below is an outline of the main steps of the genetic algorithm procedure used to select an optimal model:

Step 1) Initiate the GA by creating a *population* consisting of stochastically produced *individuals*. In our case, an individual represents a unique model specification. That is, a unique selection and grouping of variables from which the factors are created. For each individual we estimate the corresponding model and run the forecasting scheme described earlier.

Step 2) Assign each individual a *fitness value*. The fitness value allows us to compare the individuals up against each other, indicating which individuals are "fit" and which are "unfit". In our case, a measure of forecast error is used as fitness value[23], and we specify that lower measures indicate fitter individuals.

Step 3) Create a new population of individuals based on the current population using operators such as *crossovers*, *mutations*, and *selection*. The fittest individuals are transferred directly over to the new population, as these represent better forecasting models. Further, as a way of exploring the search space we create new unexplored combinations by combining individuals (e.g. select half the variables from one set-up and half the variables from another set-up, and combine them into a new set-up) or by slightly altering individuals (e.g. switching a current

[23]We use the MASE metric, which is introduced in Section 5.3.5.

variable with a new one). In our case, an example would be to pick some blocks from each of two model setups, and combine the blocks into a new setup to get a new individual. Another example is just altering a setup by adding or removing a variable.

Step 4) Repeat Step 2 and 3 until some convergence criterion is met. The fittest individuals in the final population represent models that are likely to be close to some local optimum in the search space.

Following this procedure, we have obtained a selection of models which produce the best forecasts among the models in the given search space. These models' performances are further assessed to find the very best performing model out-of-sample. The model that performs best on average is chosen as the final, *superior*, model.

5.3.4 Model Estimation

When using numerical estimation procedures, the model selection is intertwined with the estimation procedure as one has to decide a convergence criterion for the model estimation, and the models must be estimated in order to compare them. This is an important decision in estimation of forecasting models, as the models should avoid becoming *over-fitted* to the in-sample data. Over-fitting occurs when estimated models correspond too closely to a particular data set, and may therefore fail to predict future observations reliably.

We estimate the dynamic factor models using the EM algorithm as outlined in Section 3.2. The EM algorithm requires us to set the convergence threshold (τ). A lower threshold results in a greater fine-tuning of the model to the in-sample data, i.e. lowering the threshold gives an increasingly over-fitted model. We set the threshold value based on an empirical performance assessment of out-of-sample forecasts, letting the convergence threshold vary. We find empirically that $\tau = 0.001$ yields the best results.

5.3.5 Benchmark Models

We specify several benchmark models so that we can compare the DFMs' forecasting performance against widely used forecasting methods. The benchmarks included are both univariate and multivariate models often applied to forecasting. Below, we give a brief description of each model and the motivation for including them. We refer to Appendix D for a full description of each model. All models are used to forecast both Atlantic salmon and Brent Crude, and are optimally specified using the initial in-sample data. All models follow the forecasting strategy outlined in Section 5.3.1.

Univariate models

Hyndman (2010) suggests a minimum of two specific benchmarks for prediction; one naïve method and one standardized method such as an ARIMA model. ARIMA stands for autoregressive integrated moving average, and represents a group of univariate models. First, we include a naïve no-change model, formally an ARIMA$(0, 1, 0)$. We note that the naïve no-change model is applied to log price levels; that is, the model assumes that the price is constant from one period to the next, implying a zero return. The naïve no-change model is often included as a benchmark model due to its simplicity but also because it is widely recognized as a tough-to-beat benchmark.

Next, we include a more general ARIMA(p, d, q) model. The lag orders p and q are optimally specified according to the AIC.

Multivariate models

As the DFM is a multivariate model, we find it reasonable to include another multivariate model to compare performance against. Multivariate forecasting models benefit from having access to explanatory variables that may add predictive power to the models. Therefore, we include a VAR(p) model. VAR models have been extensively used in the literature to forecast commodity price changes, and is therefore a solid benchmark to include. We provide the VAR models with five key predictors recognized by Sandaker et al. (2017) and Westgaard et al. (2017) for Atlantic salmon and Brent Crude oil, respectively. Further, we specify the models' lag orders p according to the optimal AIC. The predictors included in the VAR, as well as further information about the model specification, can be found in Appendix D.

5.3.6 Performance Assessment

The performance of a given forecast can be assessed from both statistical and economic perspectives. Bloznelis (2016) categorizes statistical performance criteria into four groups; (i) measures related to the absolute forecast errors, (ii) measures of forecast accuracy relative to some naïve model for a specified sample, (iii) measures indicating whether a forecast is more accurate only in sample or also in population, (iv) measures indicating whether the forecast errors are unpredictable, i.e. whether the forecast errors possess properties similar to that of white noise. We assess the model performances using measures within all these categories. An economically relevant performance measure is how much a market participant could gain from utilizing a given forecast as compared to a benchmark forecast. Our focus will be on the statistical performance criteria, but we also illustrate the use of an economic performance assessment.

As a measure related to the absolute forecast errors we include the *root-mean-square forecast error (RMSE)*, which represents the sample standard deviation of the forecast errors. Large forecast errors are penalized heavily in this measure. Letting y_t denote the realized values and $\hat{y}_t$ denote the forecasts, RMSE is defined as follows:

$$\text{RMSE}\big(h, y_t, \hat{y}_t\big) = \sqrt{\frac{1}{t_1 - t_0 + 1} \sum_{t=t_0}^{t_1} \Big[y_{t+h} - \hat{y}_{t+h|t}\Big]^2} \tag{15}$$

We also choose to include the relative RMSE (rRMSE) metric suggested by Emiris (2016), which is just the ratio of two models' RMSE measures. This is used as an indicator of relative performance; values below unity suggests that model 1 outperforms model 2 out-of-sample and vice versa. We define rRMSE as follows:

$$\text{rRMSE}\big(h, y_t, \hat{y}_{t,1}, \hat{y}_{t,2}\big) = \frac{\text{RMSE}\big(h, y_t, \hat{y}_{t,1}\big)}{\text{RMSE}\big(h, y_t, \hat{y}_{t,2}\big)} \tag{16}$$

We include the *mean absolute scaled error (MASE)* as a measure of forecast accuracy relative to a naïve model. MASE is defined as the mean absolute error (MAE) out-of-sample of the given model relative to the *in-sample* MAE of a naïve no-change forecast. Alexander (2008b) states that the MASE metric has significant interpretation benefits; a MASE of 0.5 implies that the model's average forecast error is half of the average forecast error of a naïve no-change model in-sample.

Hyndman and Koehler (2006) define MASE as follows:

$$\text{MASE}(h, \mathbb{M}) = \frac{1}{t_1 - t_0 + 1} \sum_{t=t_0}^{t_1} \left[\frac{|e_{t+h}|}{\frac{1}{t_0} \sum_{t=2}^{t_0} |y_t - y_{t-1}|} \right], \tag{17}$$

$$\text{where} \quad e_{t+h} = y_{t+h} - \hat{y}_{t+h|t} \text{ (forecast error)}$$

We apply the widely used Diebold-Mariano (D-M) test (Diebold and Mariano, 1995; Diebold, 2015) which tests the null hypothesis that two losses due to two forecast errors are from the same population. If the null hypothesis is rejected the losses are deemed to be drawn from significantly different populations, indicating that the model that performed best is likely to also be better over time. We will let rRMSE indicate the relative performance, and complement that measure with the D-M test statistic. One must specify a loss function to use in the D-M test; as we use a squared error metric (RMSE), we use squared forecasting errors as a loss function, in line with Emiris (2016). See Appendix C for the details on the D-M test.

We perform statistical tests on the forecast residuals to assess whether their properties resemble those of white noise, which would indicate that they are unpredictable. The properties of white noise are *normality*, *insignificant serial correlation* and *homoscedastisity*. These properties are tested through the Jarque-Bera test, Ljung-Box test and Engle's test of autoregressive conditional heteroscedasticity (ARCH). All statistical tests are defined in Appendix C.

Finally, we include sample correlation and *hit rate*. The sample correlation is simply the correlation between the forecasted and realized values. *Hit rate* is defined as the percentage of forecasts that correctly predict the direction of the realized value, i.e. having the same sign (+/-). It is formally defined as:

$$\text{Hit Rate}(h, \mathbb{M}) = \frac{1}{t_1 - t_0 + 1} \sum_{t_0}^{t_1} \mathbb{I}\{sign(e_{t+h}) = sign(y_{t+h}^{obs} - y_t^{obs}) \tag{18}$$

The hit rate is neither a pure accuracy nor correlation test metric in a strict sense, but does provide complementary information regarding predictive power on a more practical economic level. In fact, the metric is easily related to risk management on a general basis, which will be explored further in the results section.

5.4 Empirical Results

This section presents the selected DFMs and provides a statistical evaluation of the 1-3 month forecasts produced. We consider Atlantic salmon and Brent Crude oil separately.

5.4.1 Atlantic Salmon

We present the DFM identified in the model selection process as the best-performing, before we assess its performance and compare it to the benchmark models and literature.

Dynamic Factor Model

The final model contains six variables, some both lagged and current; the trout price (TRO), meat index (MEA), salmon consumption in Europe (CEU), salmon consumption in emerging markets (CEM), currency pair EUR/NOK (EURNOK), and Marine Harvest Group (MHG) stock price. Table (7) reports which variables each block factor, f_j, extracts information from, along with each variable's corresponding factor loading ($\lambda_{i,j}$).

Table 7: *DFM model set-up used to forecast Atlantic salmon*

f_1	$\lambda_{i,1}$	f_2	$\lambda_{i,2}$	f_3	$\lambda_{i,3}$	f_4	$\lambda_{i,4}$	f_5	$\lambda_{i,5}$
ΔPSALM_t	0.44	ΔPSALM_t	-0.18	ΔPSALM_t	0.09	ΔPSALM_t	-0.09	ΔPSALM_t	-0.22
ΔTRO_t	0.32	ΔTRO_{t-3}	0.01	ΔTRO_{t-3}	0.02	ΔEURNOK_t	-0.36	ΔMHG_{t-1}	-0.42
ΔCEM_{t-6}	0.36	ΔCEM_t	0.04	ΔCEM_t	0.43	$\Delta\text{EURNOK}_{t-3}$	-0.21		
ΔCEU_{t-6}	0.32	ΔMEA_t	0.03	ΔCEU_t	0.55	$\Delta\text{EURNOK}_{t-6}$	-0.34		

The resulting factor representation used to predict Atlantic salmon spot price returns is:

$$\Delta\text{PSALM}_{t+h} = \lambda_1 f_{1,t+h} + \lambda_2 f_{2,t+h} + \lambda_3 f_{3,t+h} + \lambda_4 f_{4,t+h} + \lambda_5 f_{5,t+h} + \epsilon_{t+h} \quad (19)$$

The estimated latent factors are presented in Figure (8). We observe strong seasonal patterns in factors f_1 and f_3, capturing the known seasonality effects associated with the salmon market. Specifically, factor f_1 represents co-variation between Atlantic salmon, consumption in EU lagged six months, consumption in EM lagged six months, and trout. The respective factor loadings show that all variables respond similarly, both in direction and in size, to shocks in f_1. Positive

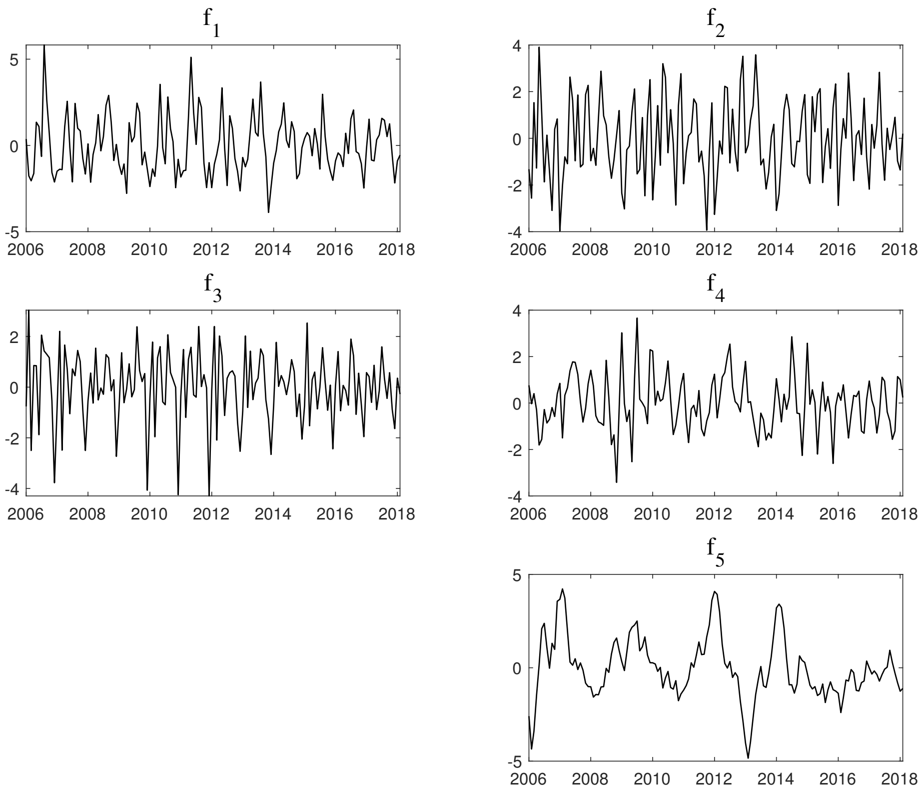

Figure 8: *Estimated latent factors for forecasting of Atlantic salmon.*

shocks implies increased consumption in both EU and EM and increased prices of both salmon and trout *six months later*. It is reasonable to think that the factor represents certain demand effects affecting the highly similar Norwegian seafood products. Factor f_5 is interesting, as a positive shock to it increases the MHG stock return this month and the salmon spot price *next month*. This could indicate price discovery effects in the MHG stock. In general however, we emphasize that making economic interpretations based on the factors and their associated variables is not straight-forward, as the factor dynamics are complex and hard to entangle. The economic interpretability is of minor interest in this application of the DFM framework.

Model Performance

The full performance assessment of the models is found in Table (8), where all forecasting accuracy measures across all models and horizons are presented.

Table 8: *Model performance metrics for the out-of-sample period from April 2015 to April 2018, Atlantic Salmon*

			Forecasting accuracy			*Forecasting errors* [a]		
Horizon	*Model*	*rRMSE* [b]	*MASE*	*Hit Rate*	*Correlation*	*JB*	*LB*	*ARCH*
$h = 1$	DFM	-	0.736	0.750	0.636	0.746	0.627	0.385
	No-change	0.820***	0.924	-	-	0.610	0.754	0.929
	ARIMA	0.745***	1.091	0.500	0.100	0.220	0.248	0.569
	VAR	0.860**	0.910	0.667	0.286	0.410	0.620	0.929
$h = 2$	DFM	-	0.762	0.750	0.617	0.777	0.516	0.445
	No-change	0.815***	0.924	-	-	0.610	0.754	0.929
	ARIMA	0.698***	1.187	0.417	−0.039	0.155	0.277	0.835
	VAR	0.824***	0.954	0.556	0.197	0.240	0.799	0.822
$h = 3$	DFM	-	0.767	0.694	0.62	0.807	0.581	0.501
	No-change	0.906***	0.924	-	-	0.610	0.754	0.929
	ARIMA	0.799***	1.138	0.417	−0.068	0.169	0.521	0.926
	VAR	0.881***	0.983	0.556	0.062	0.448	0.912	0.529

[a] *p-values reported for Jarque-Bera (JB) test for normality, Ljung-Box (LB) test for autocorrelation, Engle's ARCH test for heteroscedasticity. From a white noise perspective, large p-values indicate desirable properties.*

[b] *rRMSE is reported as DFM relative to the given model, values below unity indicate that the DFM outperforms the model. Asterisks (*, **, ***) denote that the H_0 of equal forecast accuracy between DFM and model is rejected at the (.10, .05, .01) significance level, respectively (Diebold-Mariano test).*

The broader story told by Table (8) is that the DFM produces superior forecasts across all horizons and that all measures seem to support this claim. First, all *rRMSE*[24] scores are in the range of 0.69-0.91, well below unity, stating that the DFM outperforms all benchmarks markedly. The D-M test statistic is rejected at the 0.01 significance level in all cases but one (where it is rejected at the 0.05 significance level), indicating that the DFM's forecasts are *consistently*[25] superior compared to all benchmark models. The DFM's hit rates are all greater than 0.69, even as high as 0.75 at the 1- and 2-month horizons, which are large improvements relative to flipping a coin at the beginning of each month.

[24] rRMSE is reported as DFM relative to the given model, values below unity indicate that the DFM outperforms the model.

[25] By that we mean that our model is generally better, independent of sample period.

The DFM's abovementioned performance may also be assessed visually, which is done in different ways in Figure (9). It is clear from Figure (9.I) that the model largely manages to capture the dynamics of the Atlantic salmon returns. This is further illustrated in Figure (9.III) where the price development, found by accumulating the forecasted returns, is shown. The squared forecast errors, often referred to as realized error volatility, in Figure (9.II) are consistently smaller than the other models, with the exception of the spike in March 2017. Looking at these squared errors, we see that DFM manages to capture high peaks and low downturns better than the other models, which exhibit very large squared errors at certain points. Figure (9.IV) shows that the DFM accumulates about 2/3 and 1/2 of the realized forecasting errors relative to the VAR and ARIMA, respectively. This is perhaps the most illustrative picture of the DFM's relative performance.

As the literature review uncovered, there are only two other studies that are directly comparable in that they directly forecast the price or returns of Atlantic salmon. Bloznelis (2018) considers weekly prices, but provides forecasts up to 5 weeks ahead. Of the 16 models in total tested by Bloznelis (2018), all report a *MASE* score above 1.25 at the 4 and 5 week horizons, horizons that should be comparable to our 1-month horizon. Over any horizon, his best-performing model is the vector error correcting model (VECM), which report a MASE score of 0.94 for a 1-week horizon. In comparison, our 1 month-ahead score is 0.74. This indicates that monthly data might be more suitable for forecasting Atlantic salmon returns beyond the immediate short-term. This seeming improvement on the weekly forecasts could be a result of the lower frequency data (monthly) containing less noise, but could also simply be a result of the model selection, i.e. the DFM simply performs better at forecasting than the other models. Assessing the performance of the DFM applied to weekly prices would be an interesting exercise.

The fact that Guttormsen (1999) is the only other comparable study highlights the need for more benchmark studies within this field. He predicts weekly prices of different weight classes of salmon and achieves, somewhat remarkably, hit rates well above 0.80 for most weight classes and some even above 0.90 by using a Classical Additive Decomposition (CAD) method. However, none of the models he proposes are able to produce forecasts errors that beat the naïve no-change model on a general basis. His study must be considered somewhat outdated anyway, as there have been substantial changes to the salmon market the last two decades, including considerably higher price volatility.

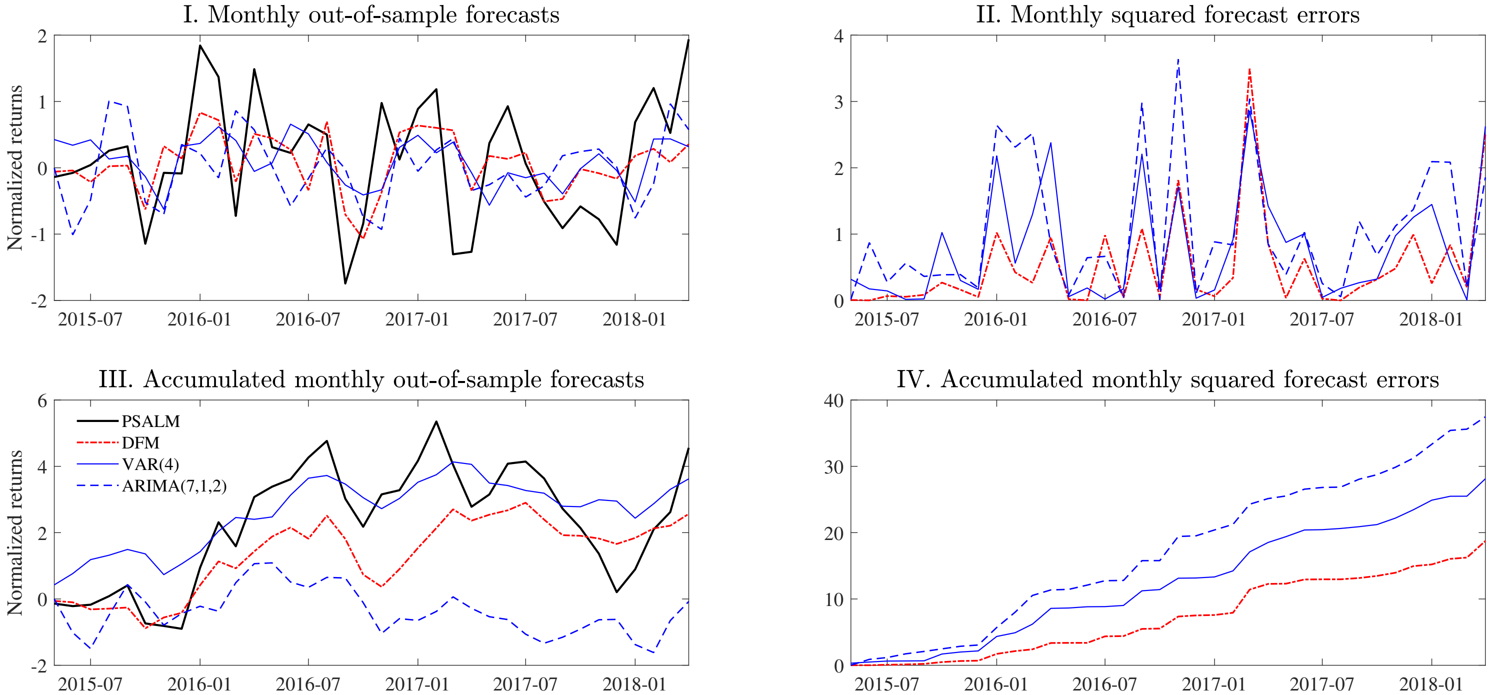

Figure 9: *Out-of-sample forecasts for 1-month-ahead Atlantic salmon returns from April 2015 to April 2018, including evaluation of forecast errors*

The forecast error diagnostics show that none of the test hypotheses can be rejected at the 10% significance level for any of the included models' forecast errors, indicating that we cannot reject that any of the forecast errors possess white-noise properties. However, these tests suffer to a certain degree from the modest sample size. Nonetheless, both the DFM and no-change model report p-values above 0.5 across all test and all horizons[26], indicating that the forecast errors to a large degree possess desirable properties. The histogram of the forecast errors for the DFM at the 1-month horizon, shown in Figure (10), shows that there are some abnormalities compared to the fitted normal distribution. This indicates that there may be more information that could be captured. However, summarizing the results from this exercise, the DFM appears to possess desirable properties and seems like a suitable model for forecasting Atlantic salmon returns.

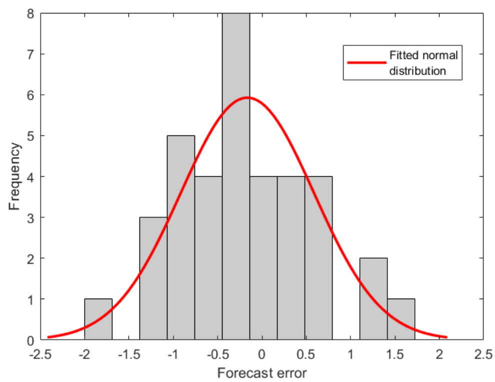

Figure 10: *Histogram of forecast errors for DFM, 1-month horizon, Atlantic salmon*

Many of the metrics discussed in this section can be considered as somewhat esoteric econometrics. *Price direction*, however, should be a fairly intuitive concept for industry participants to utilize in their planning and decision making. With a hit rate of 0.75, our model-based forecasts predict the correct direction of 1-month ahead returns 75% of the time, on average. We consider a simplified example to illustrate the potential economic value associated with this. During the out-of-sample window used in this analysis, the *average difference* in monthly salmon prices was 0.92 USD/kg. This fluctuation imposes a great deal of uncertainty onto salmon farmers on the issue of e.g. harvest-timing. Lets assume

[26]The exception being the DFM at h=1, with a p-value of 0.385 for Engle's ARCH test

that we want to plan the harvest of 1 tonnes of Atlantic salmon. The following choice must be made: harvest this month or wait. This choice will be based mainly on two pieces of information: 1) the expected price development, and 2) operational costs incurred (saved) by delaying (carrying out) harvest. Excluding operational costs from the calculations, and given the average sample price movement given above, the expected profit of using our DFM forecast in such a situation is: $1,000\text{kg} \cdot 0.92\text{USD/kg}(0.75 - (1 - 0.75)) = 460$ USD. This simple example illustrates how accurate forecasts can reduce risk and increase expected profits.

5.4.2 Brent Crude Oil

We present the DFM identified in the model selection process as the best-performing, before we assess its performance and compare it to the benchmark models and literature.

Dynamic Factor Model

The model contains six variables, some both lagged and current; PHLX oil service sector index (OSX), metal index (MET), OPEC production volumes (OPEC_PR), US net petroleum imports (IMP), trade weighted US dollar (USD) and the Kilian index (KIL). In Table (9), we show which variables each block factor, f_j, extracts information from, along with each variable's corresponding factor loading ($\lambda_{i,j}$).

Table 9: *DFM model set-up used to forecast Brent Crude oil*

f_1	$\lambda_{i,1}$	f_2	$\lambda_{i,2}$	f_3	$\lambda_{i,3}$	f_4	$\lambda_{i,4}$
ΔBCO_t	0.35	ΔBCO_t	0.12	ΔBCO_t	0.02	ΔBCO_t	0.34
ΔOSX_t	0.25	ΔOSX_t	0.19	ΔIMP_{t-3}	0.57	ΔUSD_t	-0.23
ΔMET_t	0.54	ΔOSX_{t-1}	0.13			ΔKIL_t	0.12
		$\Delta\text{OPEC\_PR}_t$	-0.63				

The resulting factor representation used to predict Brent Crude oil spot price returns is:

$$\Delta\text{BCO}_{t+h} = \lambda_1 f_{1,t+h} + \lambda_2 f_{2,t+h} + \lambda_3 f_{3,t+h} + \lambda_4 f_{4,t+h} + \epsilon_{t+h} \tag{20}$$

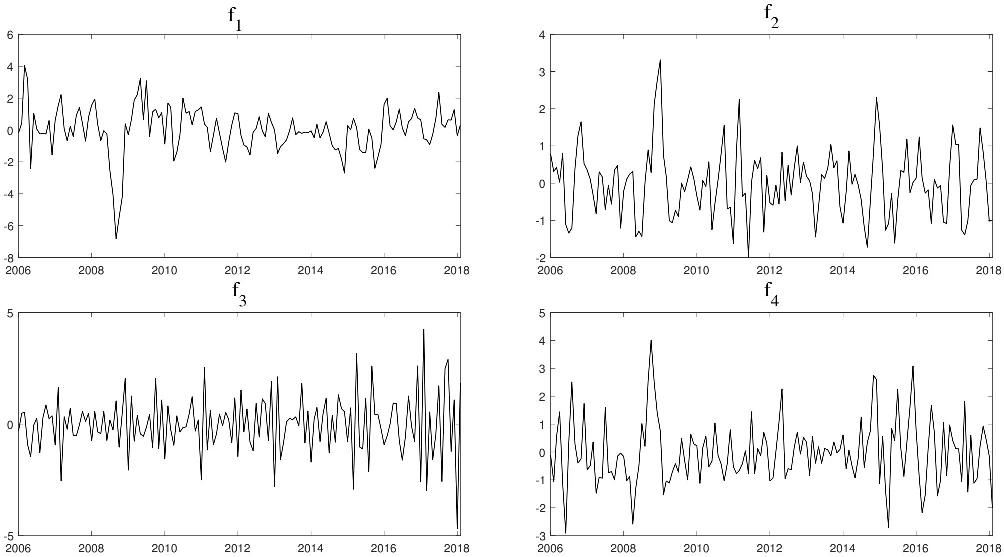

Figure 11: *The estimated factors in DFM used to forecast Brent Crude oil.*

Estimated factors are presented in Figure (11). The first factor, f_1, represents contemporaneous co-variation between oil, the OSX index and the metal index. That is, a unit shock to the factor would shift all three variables upwards, albeit with different magnitude. The factor may capture trends associated with the business cycle, supported graphically by the clearly pronounced effects of the global financial crisis. We note that in Application A, the metals sector was found to be the most co-varying with the global factor. The remaining factors do not offer similar appealing interpretations. However, the total set of latent factors provide solid oil price forecasts.

Model Performance

The full performance assessment of the models are found in Table (10), where all forecasting accuracy measures across all models and horizons are presented.

Table 10: *Model performance metrics for the out-of-sample period from April 2015 to April 2018, Brent Crude oil*

			Model Performance			*Forecast Errors*[a]		
Horizon	*Model*	*rRMSE*[b]	*MASE*	*Hit Rate*	*Correlation*	*JB*	*LB*	*ARCH*
$h = 1$	DFM	-	0.894	0.694	0.638	0.388	0.943	0.696
	No-change	0.798**	1.069	-	-	0.266	0.913	0.661
	ARIMA	0.715***	1.192	0.611	−0.061	0.970	0.847	0.087
	VAR	0.687*	1.148	0.556	−0.168	0.078	0.873	0.052
$h = 2$	DFM	-	0.955	0.667	0.491	0.741	0.736	0.456
	No-change	0.877	1.069	-	-	0.266	0.913	0.661
	ARIMA	0.752*	1.120	0.528	−0.337	0.187	0.913	0.766
	VAR	0.845*	1.038	0.667	−0.025	0.195	0.904	0.151
$h = 3$	DFM	-	0.972	0.639	0.218	0.838	0.814	0.731
	No-change	0.971	1.069	-	-	0.266	0.913	0.661
	ARIMA	0.877*	1.187	0.472	−0.215	0.404	0.854	0.395
	VAR	0.965	1.033	0.667	0.005	0.117	0.917	0.147

[a] *p-values reported for Jarque-Bera (JB) test for normality, Ljung-Box (LB) test for autocorrelation, Engle's ARCH test for heteroscedasticity. From a white noise perspective, large p-values indicate desirable properties.*

[b] *rRMSE is reported as DFM relative to the given model, values below unity indicate that the DFM outperforms the model. Asterisks (*, **, ***) denote that the H_0 of equal forecast accuracy between DFM and model is rejected at the (.10, .05, .01) significance level, respectively (Diebold-Mariano test).*

The DFM outperforms the benchmark models at all horizons, but is best at predicting 1 month ahead. At the 1-month horizon, the DFM reports 20% lower RMSE scores than any other model. At this horizon, its forecasts are found to be significantly superior to all benchmark models, with the Diebold-Mariano statistic being rejected at the 0.10 significance level or less. It performs significantly better than the ARIMA model across all horizons, which is clearly inferior to the other models at all horizons. It should be noted that ARIMA models need quite a large estimation sample in order to perform optimally. Nevertheless, our results might indicate that univariate models are not ideal for modelling Brent Crude oil returns. Further, the DFM manages to keep the MASE below unity across all horizons, a feat none of the benchmark models manage at any horizon. At the 2-month horizon, the DFM still delivers the best forecasts, although not as significantly so as it managed for 1-month ahead predictions. At the 3-month horizon, it still has a small edge on the others, yet smaller as the horizon increases. It should be noted that the hit rate stays above 60% across all horizons, even when performance is deteriorating.

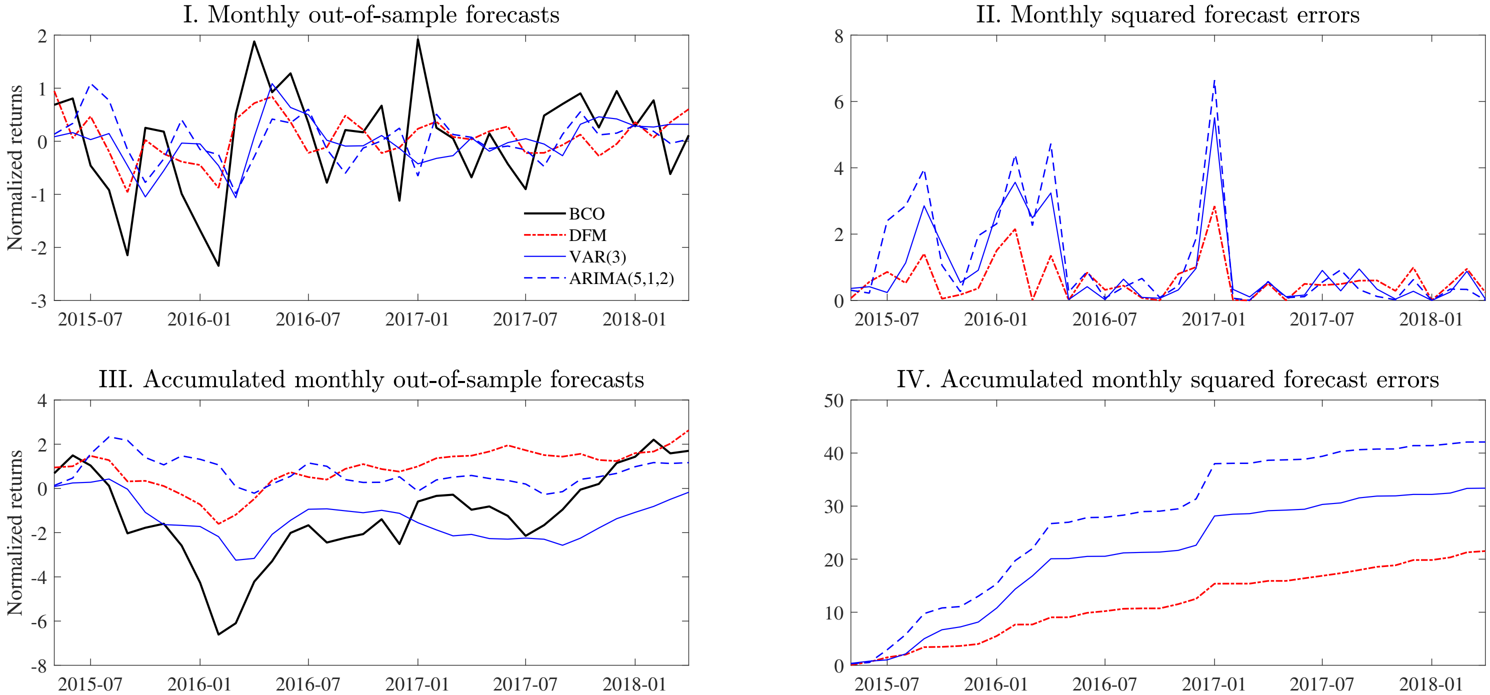

Figure 12: *Out-of-sample forecasts for 1-month-ahead Brent Crude oil returns from April 2015 to April 2018, including evaluation of forecast errors*

Examining the forecasts and realized returns in Figure (12.I), one can see that the DFM has managed to forecast large returns better than the benchmarks. This becomes clearer looking at the squared errors in Figure (12.II) and (12.IV), where we observe that the DFM manages to forecast better in volatile periods while all models achieve decent results in normal periods. Together, the figures report that the DFM accumulates significantly lower forecast errors and thereby provides more reliable indications on the future price trajectory than the other models.

Our forecasting model's performance is comparable to the very best results found in literature. Both Baumeister and Kilian (2012) and Chen (2014) report MSPE values of their models relative to a no-change model. Simply by finding the square root of these MSPE ratios, we get measures that are directly comparable to the RMSE of the DFM relative to the no-change model. At a 1-month horizon Baumeister and Kilian (2012) report a ratio of 0.905, and Chen (2014) reports a ratio of 0.848, with forecasts that are significantly superior to the no-change model (at 5% significance level according to the D-M test). Both these studies are inferior to our 1-month ahead results (0.798). Further, it is interesting to note that the oil service sector index (OSX) plays such a prominent role in the selected DFM (see Table (9)), which supports Chen (2014)'s main finding that there is substantial predictive content in oil-sensitive stock prices. We note that the models proposed by Baumeister and Kilian (2012) perform the best at forecasting 3 months ahead, where he reports an RMSE ratio of 0.862, significantly superior to the no-change model, and also outperforming the DFM. Seeing that the DFM's performance is steadily declining with horizons, it is reasonable to think that the predictors included in the model are more suited for shorter-term forecasting.[27] Further, the fact that these two models perform best at different horizons may indicate that a horizon-specific forecasting strategy could be more fitting when forecasting crude oil prices, i.e. specifying a unique model for each horizon and forecast directly not iteratively.

We find that the DFM is significantly better at predicting correct 1-month ahead price change direction relative to the abovementioned literature. Westgaard et al. (2017), who use mostly the same predictors as this study, report a hit rate of 0.58. Chen (2014) reports a hit rate of 0.62 and Baumeister and Kilian (2012) report one of 0.55 (although for WTI Crude oil). Thus, with a reported 1-month ahead hit rate of 0.69, the DFM predicts correct monthly direction about 70 % of

[27] We note that the GA-assisted model selection used 1-month MASE as a selection criterion.

the time. Models with hit rates this high should be of interest to anyone speculating in the oil market.

The forecast error diagnostics show that both the DFM and the no-change model manages quite well across all horizons, as no test hypotheses can be rejected at the 0.10 significance level. Both the ARIMA and the VAR struggle with heteroscedasticity, implying that there may be periods where they perform far worse than others. Looking at Figure (12.II) again, we see that they perform quite poorly in the period from April 2015 to April 2016 but quite well in the period of February 2017 to April 2018, causing the Engle's ARCH test null hypothesis of homoscedasticity to be rejected at the 0.10 significance level. It is worth noting that the DFM's reported p-value for the Jarque-Bera test at the 1-month horizon is slightly low, and looking at the histogram in Figure (13) the errors exhibit some non-normal features. The Brent Crude price is affected by a multitude of factors, and it is reasonable to think that there is information not currently captured by the model. However, it is clear that the model has been able to capture some information with predictive value, as it outperforms the other models.

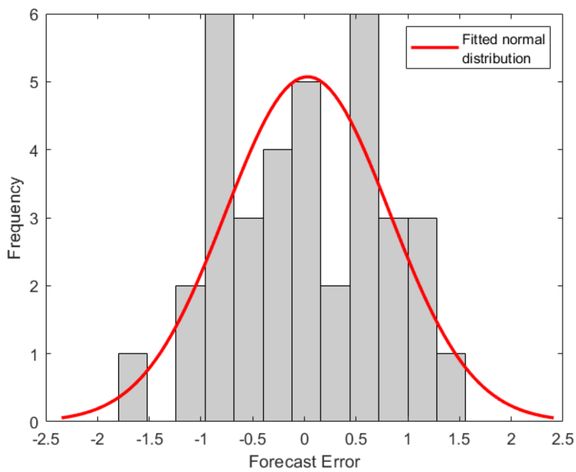

Figure 13: *Histogram of forecast errors for DFM, 1-month horizon, Brent Crude oil*

5.5 Concluding Remarks

In this study, we have utilized a DFM framework to forecast Atlantic salmon and Brent Crude oil spot price returns. The data used ranges from January 2006 to April 2018, primarily consisting of predictors recognized in the literature as having predictive content. We specified a DFM for each commodity, assisted by a genetic algorithm-based model selection procedure. We computed 1-3 month forecasts using the DFM and three other benchmark models over a 36-month out-of-sample period, following an expanding-window approach. The forecasts were assessed using a wide selection of performance measures.

We find that the DFM improves heavily upon the benchmark models in predicting Atlantic salmon returns across all horizons. It is actually found to be superior to the others at the 1% significance level. Neither Bloznelis (2018) nor Guttormsen (1999) are able to find models that consistently forecast better than the naïve no-change model. Our novel approach, on the other hand, outperforms the no-change model, with both MASE scores and RMSE ratios well below unity, and does so consistently according to D-M test statistic. This study definitely serves as a contribution to the very limited literature concerning forecasting of salmon prices. The approach proposed in this study shows great promise in predicting Atlantic salmon returns and should be of great interest to market participants.

The DFM also excels at forecasting Brent Crude returns 1 month ahead. At this horizon it is both significantly superior to the benchmark models and improves upon results found in the literature. It still improves upon the other benchmarks when forecasting longer than a month ahead, however not by as much and not significantly so. We find results in the literature seemingly outperforming the DFM at the 3-month horizon, indicating that the variables included in the DFM might contain predictive content most relevant at a shorter horizon. Adapting a direct, horizon-specific, forecasting approach may be more suitable for forecasting Brent Crude oil returns. The DFM's best results are very promising, and the framework deserves further attention.

6 Conclusion

The objective of this thesis has been to investigate the price movements of what are arguable two of the most interesting commodities in a Norwegian context, Atlantic salmon and Brent Crude oil, through two specific implementations of the general DFM framework estimated via the EM algorithm.

In *Application A*, we estimated the level of co-movement in global commodity markets, and orthogonally decomposed individual commodity returns into global, sectoral, and idiosyncratic components. This decomposition allowed us to estimate the extent to which global, sectoral and commodity-specific factors contribute to commodity price movements. The study yielded some interesting empirical results.

Focusing in on our main subjects of analysis, we find that a moderate, but significant, share of the movements in the price of both Atlantic salmon and Brent Crude over the sample period (1980-2018) can be attributed to shocks to a global factor representing common, demand-driven trends in global commodity markets. Interestingly, a sub-sample analysis reveals that, since 2000, the global factor's importance has significantly increased for both commodities (Brent Crude in particular), indicating an increased level of integration with global commodity markets. Further, we find that the DFM forms a crude, but effective, tool for gaining genuine insights into the anatomy of specific price movements for these commodities. Solely based on readily available commodity price data, the model output can assist in the interpretation of specific price events, and scrutinize existing narratives on the main drivers of these shocks.

While designing and implementing the DFM has been a demanding econometric exercise, we believe that actually making use of the model should be manageable to most. As a result, we argue that the DFM should be considered a valuable addition to the analysis toolkit used by stakeholders seeking a better understanding of commodity price composition.

Through *Application B* we have assessed the dynamic factor model's applicability to the task of forecasting monthly spot prices for Atlantic salmon and Brent Crude oil, and at the same time tested whether forecasting based on *latent* factors can provide economic value in this context. Our forecasting results suggest that the DFM performs very well in predicting the two commodities' returns on a general basis, and reports forecast errors possessing largely desirable properties (*white noise*).

The DFM outperforms all benchmarks at forecasting Atlantic salmon and Brent Crude returns across all horizons. For Atlantic salmon, we find that our forecasts are significantly superior to all benchmark forecasts at a 1% significance level, according to the D-M test. The DFM represents a solid contribution to a limited body of research concerned with directly forecasting the price of Atlantic salmon. Regarding Brent Crude, the DFM is best at forecasting 1 month ahead, where its forecast are significantly superior to all benchmark forecasts and outperforms results found in the literature. At longer horizons, it still outperforms the benchmark models, though not *significantly* so.

Our results from Application B strongly indicate that there is value to be gained from forecasting based on latent, estimated factors representing co-variation *within* a set of recognized predictor variables, rather than based on the predictor variables directly.

In conclusion, our thesis has joined the ranks of a growing number of studies demonstrating the versatility and usefulness of applying dynamic factor models to complex econometric problems.

Further Research

Since our DFM forecasting approach has yielded such promising results, a natural extension of our work would be to extend the scope of Application B. For salmon farmers, for instance, the addition of both shorter and longer term forecasts than were provided in this thesis would be of interest. The DFM could be modelled using weekly data to produce forecasts better aimed at harvesting decisions. Also, one could attempt to extend the monthly data model presented in our study in order to produce forecasts for longer horizons in order to provide decision support related to e.g. timing of smolt release.

Furthermore, the model selection problem related to our approach is inherently combinatorial, with an exhaustive search being computationally intractable. Hence, reasonable restrictions to the search space were needed to suit our computational capacity and limited time frame. Widening the search space could potentially lead to better forecasts, as this would make it possible to bring more variables into the models. In any case, we believe that our novel forecasting approach of combining a genetic algorithm for model selection with a DFM is worthy of further investigation.

References

Alexander, C., 2008a. Market Risk Analysis, Quantitative Methods in Finance. Market Risk Analysis. Wiley.
URL https://books.google.no/books?id=x3wVLNA1ofgC

Alexander, C., 2008b. Practical Financial Econometrics. Chichester, John Wiley and Sons.

Alquist, R., Coibion, O., 2014. Commodity-price comovement and global economic activity. Tech. rep., National Bureau of Economic Research.

Asche, F., Bjorndal, T., 2011. The economics of salmon aquaculture. Vol. 10. John Wiley & Sons.

Asche, F., Guttormsen, A., Roll, K. H., Tveteraas, R., 2013. Produksjonsvekst, innovasjon og miljøutfordringer. Veksten i norsk lakseoppdrett. Samfunnsøkonomen 4, 82–91.

Asche, F., Oglend, A., 2016. The relationship between input-factor and output prices in commodity industries: The case of Norwegian salmon aquaculture. Journal of Commodity Markets 1 (1), 35–47.

Bai, J., 2003. Inferential theory for factor models of large dimensions. Econometrica 71 (1), 135–171.
URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-0262.00392

Bai, J., Ng, S., 2007. Determining the number of primitive shocks in factor models. Journal of Business & Economic Statistics 25 (1), 52–60.
URL https://doi.org/10.1198/073500106000000413

Bańbura, M., Giannone, D., Reichlin, L., 2011. Nowcasting. In: The Oxford Handbook of Economic Forecasting. Oxford University Press, pp. 1–37.

Bańbura, M., Modugno, M., 2014. Maximum likelihood estimation of factor models on datasets with arbitrary pattern of missing data. Journal of Applied Econometrics 29 (1), 133–160.

Barhoumi, K., DarnÃĩ, O., Ferrara, L., 2013. Dynamic Factor Models: A review of the Literature . Working papers 430, Banque de France.
URL https://ideas.repec.org/p/bfr/banfra/430.html

Baumeister, C., Kilian, L., 2012. Real-time forecasts of the real price of oil. Journal of Business & Economic Statistics 30 (2), 326–336.

Baumeister, C., Kilian, L., 2014. Do oil price increases cause higher food prices? Economic Policy 29 (80), 691–747.
URL http://dx.doi.org/10.1111/1468-0327.12039

Baumeister, C., Kilian, L., 2016. Understanding the decline in the price of oil since June 2014. Journal of the Association of Environmental and Resource Economists 3 (1), 131–158.

Beckers, B., 2015. Forecasting the Nominal Brent Oil Price with VARs - One Model Fits All? International Monetary Fund.

Bilgin, D., Ellwanger, R., 2017. A dynamic factor model for commodity prices. Bank of Canada Staff Analytical Note 2017 (12).

Bjørnland, H., Thorsrud, L., 2014. What is the effect of an oil price decrease on the Norwegian economy. Oslo: Norges Bank.

Bjørnland, H. C., 1997. Estimating core inflation: the role of oil price shocks and imported inflation. Tech. Rep. 200, Statistics Norway, Research Department, East Lansing, Michigan. URL http://hdl.handle.net/11250/180647

Bjørnland, H. C., Thorsrud, L. A., 2016. Boom or gloom? Examining the Dutch Disease in two-speed economies. The Economic Journal 126 (598), 2219–2256.

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

Bloznelis, D., 2018. Short-term salmon price forecasting. Journal of Forecasting 37 (2), 151–169.

Bork, L., 2009. Estimating US monetary policy shocks using a factor-augmented vector autoregression: An EM algorithm approach. Creates research papers, Department of Economics and Business Economics, Aarhus University.
URL https://EconPapers.repec.org/RePEc:aah:create:2009-11

Brækkan, E. H., Thyholdt, S. B., 2014. The bumpy road of demand growth âĂŤ an application to Atlantic salmon. Marine Resource Economics 29 (4), 339–350.

Breitung, J., Pigorsch, U., 2012. A canonical correlation approach for selecting the number of dynamic factors. Oxford Bulletin of Economics and Statistics 75 (1), 23–36.
URL https://onlinelibrary.wiley.com/doi/abs/10.1111/obes.12003

Brooks, C., 2014. Introductory econometrics for finance. Cambridge university press.

Byrne, J. P., Fazio, G., Fiess, N., 2013. Primary commodity prices: Co-movements, common factors and fundamentals. Journal of Development Economics 101, 16–26.

Chen, S., 2014. Forecasting crude oil price movements with oilâĂŘsensitive stocks. Economic Inquiry 52 (2), 830–844.
URL https://onlinelibrary.wiley.com/doi/abs/10.1111/ecin.12053

Cheng, I.-H., Xiong, W., 2014. Financialization of commodity markets. Annu. Rev. Financ. Econ. 6 (1), 419–441.

De Mol, C., Giannone, D., Reichlin, L., 2008. Forecasting using a large number of predictors: Is Bayesian shrinkage a valid alternative to principal components? Journal of Econometrics 146 (2), 318–328.

Delle Chiaie, S., Ferrara, L., Giannone, D., 2017. Common factors of commodity prices. La Défense-EconomiX.

Dempster, A. P., Laird, N. M., Rubin, D. B., 1977. Maximum likelihood from incomplete data via the EM algorithm. Journal of the Royal Statistical Society. Series B (Methodological) 39 (1), 1–38.
URL http://www.jstor.org/stable/2984875

Diebold, F. X., 2015. Comparing predictive accuracy, twenty years later: A personal perspective on the use and abuse of dieboldâĂŞmariano tests. Journal of Business & Economic Statistics 33 (1), 1–1.
URL https://doi.org/10.1080/07350015.2014.983236

Diebold, F. X., Mariano, R. S., 1995. Comparing predictive accuracy. Journal of Business & Economic Statistics 13 (3), 253–263.
URL http://www.jstor.org/stable/1392185

Doz, C., Giannone, D., Reichlin, L., 2011. A two-step estimator for large approximate dynamic factor models based on kalman filtering. Journal of Econometrics 164 (1), 188–205.
URL https://EconPapers.repec.org/RePEc:eee:econom:v:164:y:2011:i:1:p:188-205

Doz, C., Giannone, D., Reichlin, L., 2012. A quasi-maximum likelihood approach for large, approximate dynamic factor models. Review of economics and statistics 94 (4), 1014–1024.

Durbin, J., Koopman, S. J., 2012. Time series analysis by state space methods. Vol. 38. OUP Oxford.

Emiris, M., 2016. A dynamic factor model for forecasting house prices in Belgium. Tech. rep., Working Paper Research.

FAO, 2016. The state of world fisheries and aquaculture - 2016 (SOFIA). Rome, Italy: Food and Agriculture Organization.

Fattouh, B., 2007. Opec pricing power : the need for a new perspective. Tech. rep., Oxford Institute for Energy Studies.

Forni, M., Giannone, D., Lippi, M., Reichlin, L., 2009. Opening the black box: Structural factor models with large cross sections. Econometric Theory 25 (5), 1319âĂŞ1347.

Forni, M., Hallin, M., Lippi, M., Reichlin, L., 2000. The generalized dynamic-factor model: Identification and estimation. Review of Economics and statistics 82 (4), 540–554.

Forni, M., Hallin, M., Lippi, M., Reichlin, L., 2005. The generalized dynamic factor model: one-sided estimation and forecasting. Journal of the American Statistical Association 100 (471), 830–840.

Geweke, J. F., 1977. The dynamic factor analysis of economic time series models. In: Aigner, D., Goldberger, A. (Eds.), Latent variables in socio-economic models. North-Holland Pub. Co., pp. 365–383.

Giannone, D., Reichlin, L., Small, D., 2008. Nowcasting: The real-time informational content of macroeconomic data. Journal of Monetary Economics 55 (4), 665 – 676.
URL http://www.sciencedirect.com/science/article/pii/S0304393208000652

Globefish, August 2006. Commodity statistics update: Salmon. Tech. rep., GLOBEFISH.

Guttormsen, A. G., 1999. Forecasting weekly salmon prices: risk management in fish farming. Aquaculture Economics & Management 3 (2), 159–166.

Hagen, R., 1994. How is the international price of a particular crude determined? OPEC Energy Review 18 (1), 127–135.

Hallin, M., Liška, R., 2007. Determining the number of factors in the general dynamic factor model. Journal of the American Statistical Association 102 (478), 603–617.
URL http://www.jstor.org/stable/27639890

Hamilton, J. D., November 2008. Understanding crude oil prices. Working Paper 14492, National Bureau of Economic Research.
URL http://www.nber.org/papers/w14492

Harvey, D., Leybourne, S., Newbold, P., 1997. Testing the equality of prediction mean squared errors. International Journal of forecasting 13 (2), 281–291.

Hoover, K. D., Perez, S. J., 1999. Data mining reconsidered: encompassing and the general-to-specific approach to specification search. The econometrics journal 2 (2), 167–191.

Hyndman, R., 2010. Benchmarks for forecasting.
URL https://robjhyndman.com/hyndsight/benchmarks/#disqus_thread

Hyndman, R. J., Koehler, A. B., 2006. Another look at measures of forecast accuracy. International journal of forecasting 22 (4), 679–688.

International Energy Agency, 2017. Key world energy statistics.

Jansen, W. J., Jin, X., de Winter, J. M., 2016. Forecasting and nowcasting real gdp: Comparing statistical models and subjective forecasts. International Journal of Forecasting 32 (2), 411 – 436.
URL http://www.sciencedirect.com/science/article/pii/S0169207015001119

Juvenal, L., Petrella, I., 2015. Speculation in the oil market. Journal of Applied Econometrics 30 (4), 621–649.

Karlsson, S., 2013. Forecasting with Bayesian vector autoregression. In: Handbook of economic forecasting. Vol. 2. Elsevier, pp. 791–897.

Kilian, L., 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. American Economic Review 99 (3), 1053–69.

Kilian, L., Murphy, D. P., 2014. The role of inventories and speculative trading in the global market for crude oil. Journal of Applied Econometrics 29 (3), 454–478.

Krolzig, H.-M., Hendry, D. F., 2001. Computer automation of general-to-specific model selection procedures. Journal of Economic Dynamics and Control 25 (6-7), 831–866.

Marcellino, M., Stock, J. H., Watson, M. W., 2006. A comparison of direct and iterated multi-step AR methods for forecasting macroeconomic time series. Journal of econometrics 135 (1-2), 499–526.

Marine Harvest, 2017. Salmon farming industry handbook 2017. Tech. rep., Marine Harvest ASA.

Miranda-Agrippino, S., Rey, H., 2015. World asset markets and the global financial cycle. Tech. rep., National Bureau of Economic Research.

Mitchell, D., July 2008. A note on rising food prices. Working Paper WPS4682, World Bank. URL http://documents.worldbank.org/curated/en/229961468140943023/A-note-on-rising-food-prices

Mitchell, M., 1996. An Introduction to Genetic Algorithms. MIT Press, Cambridge, MA, USA.

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

Ohashi, K., Okimoto, T., 2016. Increasing trends in the excess comovement of commodity prices. Journal of Commodity Markets 1 (1), 48–64.

Olimb, M., Odegaard, T. M., 2010. Understanding the factors behind crude oil price changes.

Plourde, A., Watkins, G., 1994. How volatile are crude oil prices? OPEC Energy Review 18 (4), 431–444.

Rauch, H. E., Striebel, C., Tung, F., 1965. Maximum likelihood estimates of linear dynamic systems. AIAA journal 3 (8), 1445–1450.

Sandaker, K., Mjaugeto, P. O. W., Steinshamn, K. B., 2017. Predicting the distribution of the atlantic salmon spot price using quantile regression. Master's thesis, NTNU - Norwegian University of Science and Technology.

Sargent, T. J., Sims, C. A., 1977. Business cycle modeling without pretending to have too much a priori economic theory. Working Papers 55, Federal Reserve Bank of Minneapolis.
URL https://ideas.repec.org/p/fip/fedmwp/55.html

Shumway, R. H., Stoffer, D. S., 1982. An approach to time series smoothing and forecasting using the em algorithm. Journal of Time Series Analysis 3 (4), 253–264.
URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9892.1982.tb00349.x

Solibakke, P. B., 2012. Scientific stochastic volatility models for the salmon forward market: forecasting (un-) conditional moments. Aquaculture Economics & Management 16 (3), 222–249.

SSB, 2018. Utenrikshandel med varer. https://www.ssb.no/utenriksokonomi/statistikker/muh/aar, accessed: 2018-04-13.

Stevens, P., 1995. The determination of oil prices 1945-1995: a diagrammatic interpretation. Energy Policy 23 (10), 861–870.

Tang, K., Wang, C., Wang, S., 2014. China's imported inflation and global commodity prices. Emerging Markets Finance and Trade 50 (3), 162–177.

Terazono, E., 2016. Chile's salmon sector hit by toxic algae bloom.
URL https://www.ft.com/content/1978a158-12ac-11e6-91da-096d89bd2173

Wall Street Journal, 2018. Oil's biggest benchmark, brent, needs more oil.
URL https://www.wsj.com/articles/oils-biggest-benchmark-brent-needs-more-oil-1525086000

Watson, M. W., Engle, R. F., 1983. Alternative algorithms for the estimation of dynamic factor, mimic and varying coefficient regression models. Journal of Econometrics 23 (3), 385 – 400.
URL http://www.sciencedirect.com/science/article/pii/030440768390066o

West, K. D., Wong, K.-F., 2014. A factor model for co-movements of commodity prices. Journal of International Money and Finance 42, 289–309.

Westgaard, S., Osmundsen, P., Stenslet, D., Ringheim, J., 2017. Modeling superior predictors for crude oil prices. Journal of Energy Markets 10 (2), 1–22.
URL https://ssrn.com/abstract=2983716

Wu, L. S.-Y., Pai, J. S., Hosking, J., June 1996. An algorithm for estimating parameters of state-space models. Statistics & Probability Letters 28 (2), 99–106.
URL https://ideas.repec.org/a/eee/stapro/v28y1996i2p99-106.html

Yin, L., Han, L., 2015. Co-movements in commodity prices: Global, sectoral and commodity-specific factors. Economics Letters 126, 96–100.

Appendix

A Factor plots and loadings

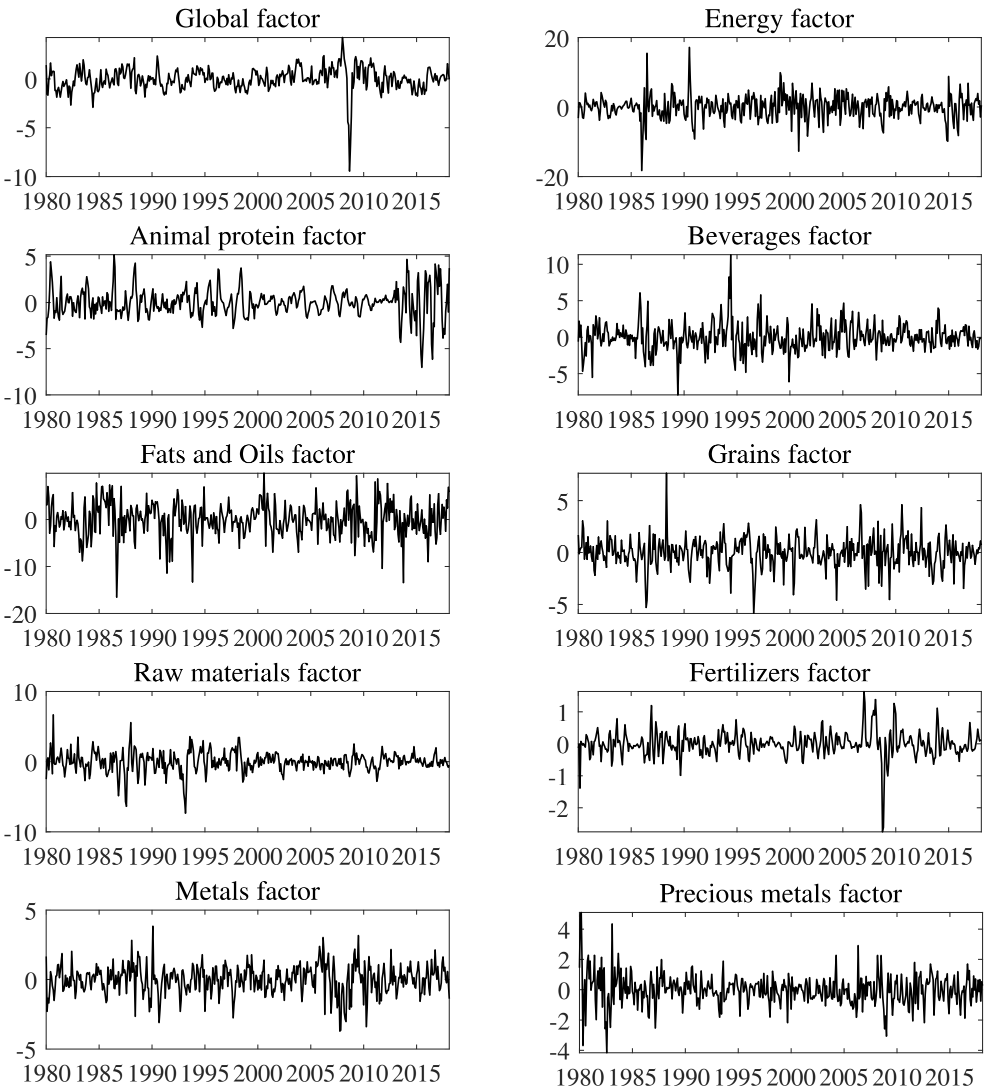

Figure 14: *Return plots of the estimated dynamic factors*

Table 11: *Estimated Model Parameters of all Commodity Log Returns.*

	Factor Loadings ($\hat{\Lambda}$)									
Variables	Global	Animal Proteins	Beverages	Energy	Fats & Oils	Fertilizers	Grains	Metals	Precious metals	Raw materials
Atlantic salmon	*0.256*	*0.063*	-	-	-	-	-	-	-	-
Beef	0.250	0.054	-	-	-	-	-	-	-	-
Chicken	0.049	0.092	-	-	-	-	-	-	-	-
Sheep	0.217	0.121	-	-	-	-	-	-	-	-
Shrimps	0.020	0.218	-	-	-	-	-	-	-	-
Cocoa	0.184	-	0.168	-	-	-	-	-	-	-
Coffee, Arabica	0.123	-	0.336	-	-	-	-	-	-	-
Coffee, Robusta	0.163	-	0.395	-	-	-	-	-	-	-
Tea	0.202	-	0.112	-	-	-	-	-	-	-
Coal, Australian	0.388	-	-	0.061	-	-	-	-	-	-
Coal, Colombian	0.403	-	-	0.053	-	-	-	-	-	-
Crude oil, Brent	*0.319*	-	-	*0.168*	-	-	-	-	-	-
Crude oil, WTI	0.316	-	-	0.155	-	-	-	-	-	-
LNG, Japan	0.106	-	-	-0.037	-	-	-	-	-	-
Natural gas, Europe	0.117	-	-	0.074	-	-	-	-	-	-
Coconut oil	0.190	-	-	-	0.191	-	-	-	-	-
Copra	0.177	-	-	-	0.135	-	-	-	-	-
Fish meal	0.309	-	-	-	0.037	-	-	-	-	-
Groundnut oil	0.252	-	-	-	0.099	-	-	-	-	-
Groundnuts	0.231	-	-	-	0.088	-	-	-	-	-
Palm oil	0.292	-	-	-	0.125	-	-	-	-	-
Soybean meal	0.343	-	-	-	0.07	-	-	-	-	-
Soybeans	0.374	-	-	-	0.077	-	-	-	-	-
DAP	0.075	-	-	-	-	0.541	-	-	-	-
Urea	0.296	-	-	-	-	0.299	-	-	-	-
Barley	0.303	-	-	-	-	-	0.266	-	-	-
Maize	0.259	-	-	-	-	-	0.577	-	-	-
Rice	0.282	-	-	-	-	-	0.013	-	-	-
Sorghum	0.258	-	-	-	-	-	0.531	-	-	-
Wheat	0.182	-	-	-	-	-	0.297	-	-	-
Aluminum	0.400	-	-	-	-	-	-	0.351	-	-
Copper	0.454	-	-	-	-	-	-	0.431	-	-
Lead	0.285	-	-	-	-	-	-	0.415	-	-
Nickel	0.316	-	-	-	-	-	-	0.402	-	-
Tin	0.379	-	-	-	-	-	-	0.191	-	-
Zinc	0.296	-	-	-	-	-	-	0.531	-	-
Gold	0.278	-	-	-	-	-	-	-	0.657	-
Platinum	0.411	-	-	-	-	-	-	-	0.531	-
Silver	0.300	-	-	-	-	-	-	-	0.669	-
Logs, Malaysia	0.040	-	-	-	-	-	-	-	-	0.657
Logs, West Africa	0.339	-	-	-	-	-	-	-	-	0.147
Sawnwood	0.113	-	-	-	-	-	-	-	-	0.392
Plywood	0.064	-	-	-	-	-	-	-	-	0.452
Woodpulp	0.268	-	-	-	-	-	-	-	-	0.139

[a] Notes: The table reports the estimated factor loadings ($\hat{\Lambda}$) for all commodities in the information set $\mathcal{I}*$ over the sample period 1980:M1 - 2018:M3. The loadings quantify the sensitivity of each commodity i to shocks in the global and sectoral factors. Note that it is the *product* of the loading and the factor that determine the respective components' contribution to the individual commodity fluctuations.

B Information Sets

Table 12: *Forecasting information set for Atlantic salmon with descriptive statistics.*

Time series, (unit)	Source	Lag(s)[a]	Trans.	*Time window*	$\bar{y}$	σ	JB[b]	LB[c]	ADF[d]
Consumption of salmon, Europe (tonnes)	Marine Harvest Group	3,6,12	Δln	YoY	0.05	0.08	0.41	251.34	-5.26
Consumption of salmon, Emerging Markets (tonnes)	Marine Harvest Group	3,6,12	Δln	YoY	0.1	0.13	44.27	298.85	-3.48
Currency pair, Euro to Norwegian krone (EUR/NOK)	Oanda.com	3,6,12	Δln	MoM	0.00	0.02	37.95	43.43	-9.48
Currency pair, Chilean dollar to US dollar (CLD/USD)	Oanda.com	3,6,12	Δln	MoM	0.00	0.03	191.05	37.39	-11.19
Norwegian trout, spot price (USD/kg)	Norwegian Seafood Council	1-6	Δln	MoM	0.00	0.07	44.70	72.74	-9.1
Norwegian mackerel, spot price (USD/kg)	Norwegian Seafood Council	1-6	Δln	MoM	0.00	0.09	43.35	37.57	-20.39
Shrimp (Mexico), spot price (USD/pound)	Worldbank.com	1-6	Δln	MoM	0.00	0.05	97.27	113.57	-5.17
Meat price index	FAO.org	1-6	Δln	MoM	0.00	0.03	25.48	171.94	-6.69
Beef (US), spot price (USD/pound)	FAO.org	1-6	Δln	MoM	0.00	0.04	67.16	70.99	-8.69
Chicken (US), spot price (USD/pound)	Worldbank.com	1-6	Δln	MoM	0.00	0.03	54.45	69.21	-11.78
Feed consumption, Norway (tonnes)	Fiskedir.no	1-6	Δln	YoY	0.00	0.12	15.43	365.97	-3.69
Sea lice occurence, Norway (#lice/fish)	Lusedata.no	3,6,9,12,15	Δln	YoY	0.03	0.45	28.22	414.12	-3.53
Sea temperature. Norway ($C°$)	Lusedata.no	3,6,9,12,15	Δln	YoY	0.00	0.13	17.76	387.35	-2.76
Smolt release, Norway (#individuals)	Fiskedir.no	9,12,15,17	Δln	YoY	0.04	0.54	696.42	15.52	-11.12
Standing biomass salmon, Norway (tonnes)	Fiskedir.no	3,6,9,12,15	Δln	YoY	0.08	0.12	629.48	219.76	-0.13
Harvest volumes salmon, Norway (tonnes)	Fiskedir.no	3,6,9,12,15	Δln	YoY	0.07	0.15	234.46	195.76	-4.67
Harvest volumes trout, Norway (tonnes)	Fiskedir.no	3,6,9,12,15	Δln	YoY	0.02	0.19	98.98	145.67	-3.22
Fishmeal, spot price (USD/metric ton)	Worldbank.com	1-6	Δln	MoM	0.00	0.04	43.8	55.17	-7.46
Rapeseed oil, spot price (USD/metric ton)	Worldbank.com	1-6	Δln	MoM	0.00	0.05	23.61	72.95	-7.35
Soybeans, US, spot price (USD/metric ton)	Worldbank.com	1-6	Δln	MoM	0.00	0.06	35.39	43.87	-8.47
Wheat, US, spot price (USD/metric ton)	Worldbank.com	1-6	Δln	MoM	0.00	0.08	18.98	30.14	-9.4
Norwegian farmed Atlantic Salmon, PSALM (USD/kg)	SSB.no	1-6	Δln	MoM	0.01	0.15	13.18	48.27	-9.02
OSLO Seafood Index, OSLOFX (USD)	Yahoofinance.com	1-6	Δln	MoM	0.05	0.08	23.05	70.46	-11.22
Marine Harvest Group (MHG) stock price (USD/share)	Yahoofinance.com	1-6	Δln	MoM	0.01	0.07	25.04	67.47	-8.46
Europe Brent Crude oil, BCO (USD/barrel)	EIA.gov	1-6	Δln	MoM	0.00	0.09	44.77	40.13	-8.38

[a] The specific lag structures are based on the variable selection findings in Sandaker et al. (2017).
[b] Critical values JB : $\chi^2_{2,0.1}$>4.61, $\chi^2_{2,0.05}$>5.99 & $\chi^2_{2,0.01}$>9.21
[c] Critical values LB : $\chi^2_{6,0.1}$>10.64, $\chi^2_{6,0.05}$>12.59 & $\chi^4_{6,0.01}$>16.81
[d] Critical values ADF : $\tau_{0.1}$<-1.62, $\tau_{0.05}$<-1.95 & $\tau_{0.01}$<-2.58

Table 13: *Forecasting information set for Brent Crude oil with descriptive statistics.*

Time series, (unit)	*Source*	*Lag(s)*	Trans.	Time window	$\bar{y}$	σ	JB [a]	LB [b]	ADF [c]
US unemployment rate (% seas. adj.)	Reuters.com	1-6	-	-	0.00	0.03	5.07	138.3	-9.82
US money supply (USD)	Reuters.com	1-6	Δln	MoM	0.01	0.01	293.27	47.89	-8.79
GDP index (USD)	Quandl.com	1-6	Δln	MoM	0.00	0	342.52	207.25	-2.87
USD index, trade weighted (#)	Reuters.com	1-6	Δln	MoM	0.00	0.02	7.18	54	-8.45
Treasury bills, 3 month yield (%)	Reuters.com	1-6	-	-	1.00	1.61	75.39	1369.42	-2.63
Agricultural Commodity Index (USD)	Worldbank.com	1-6	Δln	MoM	0.00	0.03	193.81	67.66	-7.55
Metal index (USD)	Worldbank.com	1-6	Δln	MoM	0.00	0.06	148.5	50.65	-7.65
Gold, spot price (USD)	Worldbank.com	1-6	Δln	MoM	0.01	0.03	98.2	54.34	-7.31
US Natural gas, spot price (USD)	Worldbank.com	1-6	Δln	MoM	0.00	0.04	123.41	62.66	-12.85
US petroleum imports (USD)	Worldbank.com	1-6	Δln	MoM	0.00	0.02	143.76	33.66	-5.12
Kilian index, dev. from trend (%)	personal.umich.edu/ lkilian	1-6	-	-	3.43	36.06	7.16	972.34	-2.25
OPEC production (turnover by volume)	Reuters.com	1-6	Δln	MoM	0.00	0.01	14.15	25.53	-10.13
OPEC surplus capacity (oil volume)	Reuters.com	1-6	Δln	MoM	0.00	0.11	25.61	34.7	-8.6
Non-OPEC production. (turnover by volume)	Reuters.com	1-6	Δln	MoM	0.00	0.01	2.48	85.5	-12.06
Rotary oil rigs (# in operation. US)	Reuters.com	1-6	Δln	MoM	0.01	0.07	65.56	133.66	-5.21
OECD inventories (inventory volume)	Reuters.com	1-6	Δln	MoM	0.00	0.01	0.76	61.72	-12.85
OPEC inventories (inventory volume)	EIA.gov	1-6	Δln	MoM	0.00	0.02	0.56	41.72	-18.41
US inventories (inventory volume)	EIA.gov	1-6	Δln	MoM	0.01	0.02	0.67	35.72	-9.25
Refinery utilization (% operable capacity, US)	Reuters.com	1-3	-	-	0.00	0.03	45.74	56.25	-12.57
Europe Brent Crude oil, BCO (USD/barrel)	EIA.gov	1-6	Δln	MoM	0.00	0.09	44.77	40.13	-8.38
PHLX Oil Service Sector index, OSX (USD)	Yahoofinance.com	1-6	Δln	MoM	0.01	0.04	0.57	15.94	-10.8
CBOE Volatility Index, VIX (USD)	Yahoofinance.com	1-6	Δln	MoM	0.00	0.21	19.79	32.54	-14.64
Futures spread, 6month v 1month NYMEX (%)	Quandl.com	1-6	-	-	1.84	7.15	385.42	32.24	-11.51
BB spread, ICE BofAML US High Yield (%)	Quandl.com	1-6	-	-	4.18	2.14	377.65	671.32	-0.79
NYMEX Crack Spread (USD)	Reuters.com	1-6	Δln	MoM	0.00	0.09	965.33	18.21	-10.66

[a] Critical values JB : $\chi^2_{2,0.1}$>4.61, $\chi^2_{2,0.05}$>5.99 & $\chi^2_{2,0.01}$>9.21

[b] Critical values LB : $\chi^2_{6,0.1}$>10.64, $\chi^2_{6,0.05}$>12.59 & $\chi^4_{6,0.01}$>16.81

[c] Critical values ADF : $\tau_{0.1}$<-1.62, $\tau_{0.05}$<-1.95 & $\tau_{0.01}$<-2.58

C Statistical Tests

Diebold-Mariano (1995) Statistic

The test assesses whether the loss due to forecast errors is equally great for two competing forecasting models. That is; given two plausible models, the null hypothesis (H_0) of the D-M test statistic is given by $\mathbb{E}[D_t]$=0 for all $t \in [t_0, t_1]$, where D_t=f$\big(e_{i,t}\big)$-f$\big(e_{j,t}\big)$ is the loss differential from model i vs. model j, and $f(.)$ is a specified loss function.[28] The test statistic is defined as follows,

$$\mathcal{S}_{i,j} = \frac{\overline{D}_t}{\sqrt{\hat{\sigma}/T}}, \qquad \mathcal{S} \sim N(0,1) \tag{21}$$

where $\overline{D}_t$ is the sample mean of the loss differential $D_{t_0}, \ldots, D_{t_1}$. $\hat{\sigma}$ is a consistent estimator of the asymptotic variance of $\sqrt{T\overline{D}}$. Further, we use a modified version of the statistic proposed by Harvey et al. (1997) to account for probable auto-correlation between forecast errors. This comes in handy in cases where the forecast errors exhibit weakly white noise properties.

Jarque-Bera Test

The Jarque-Bera(JB) test can be applied to any random variable whenever we need to justify an assumption of normality (see Alexander, 2008a, p.158). The JB statistic is given by:

$$\text{JB} = \frac{T-k}{6}\left(\hat{\tau}^2 + \frac{1}{4}(\hat{\kappa}-3)^2\right) \quad \sim \quad \chi^2_2 \tag{22}$$

where T is the number of observations and k is the number of regressors. The H_0 hypothesis is that the time series is equal to that of the normal distribution, while H_a hypothesis states that the series is not normal. From Equation (22), the series will be normally distributed if both skewness ($\hat{\tau}$) and kurtosis ($\hat{\kappa}$) is close to zero. The JB test rejects H_0 if the test statistic exceeds the critical value for a given significance level (α).

[28] We perform the test on normalized log returns and thus have no clear intuitive loss function. Therefore, we follow the standard approach and use absolute squared function as f(.)

Engle's ARCH test

A time series that exhibits conditional heteroscedasticity - or autocorrelation in the squared series - is said to have autoregressive conditional heteroscedastic (ARCH) properties (see Brooks, 2014, p.389). Engle's ARCH test evaluates each time series for such properties. In this study, we use the test to investigate model residuals. The null and alternative hypothesis are the following,

$$\begin{aligned} H_a : e_t^2 = \alpha_0 + \alpha_1 e_{t-1}^2 + ... + \alpha_m e_{t-m}^2 + u_t \\ H_0 : a_0 = a_1 + ... + a_m = 0. \end{aligned} \tag{23}$$

where e_t is the model residual and u_t is a white noise process. The test criteria stated in Equation (23) simply evaluate if the autoregressive constants α are statistically different from zero or not. Thus, the test rejects the null hypothesis (H_0) if the model residual exhibits ARCH properties and vice versa.

Ljung-Box Q-test

The Ljung-Box Q-test evaluates each time series for auto-correlation. Its advantage over other similar tests, like the *Box-Pierce test*, is that it tests for randomness based on a specific number of lags instead of testing for randomness for each lag. The statistic is given by

$$Q_{LB} = n(n+2) \sum_{k=1}^{m} \frac{\hat{r}_k^2}{n-k} \quad \sim \quad \chi_h^2 \tag{24}$$

where $\hat{r}_k$ is the estimated autocorrelation of the series at lag k, and m is the number of lags being tested and n is the number of observations. The H_0 hypothesis is that the time series consist of random observations, not affect by significant auto-correlation. The Ljung-Box test rejects H_0 if the test statistic exceeds the critical value for a given significance level (α).

Augmented Dickey-Fuller (ADF) Test

The ADF test is an improved Dicky-Fuller (DF) test used to investigate stationarity in time series. The advantage of the ADF test over the simple DF test is that by including necessary lagged dependent variables, it solves the problem of autocorrelation in the residuals. The test of order q, or ADF(q) test, is based on the auxiliary regression (see Alexander, 2008a, p.217):

$$\Delta X_t = \alpha + \beta X_{t-1} + \gamma_1 \Delta X_{t-1} ... + \gamma_q \Delta X_{t-q} + \epsilon_t \tag{25}$$

There exists several approaches to investigate the H_0 hypothesis. One approach is to use the usual t-statistics of the estimated coefficient $\hat{\beta}$. The critical values are dependent on the number of lags, q. In this paper, all the ADF tests are calculated with a single lag, i.e. q = 1, as the tests do not yield notably different statistics for longer lags.

D Benchmark Models

D.1 VAR(p)

Let $\mathbf{Y}_t = (y_{1t}, ..., y_{1n})^T$ denote a vector of time series variables. Then, the basic p-lag vector autoregressive VAR(p) model is given by

$$\mathbf{Y}_t = \mathbf{c} + \Pi_1 \mathbf{Y}_{t-1} + ... + \Pi_p \mathbf{Y}_{p-1} + \epsilon_t \tag{26}$$

where Π_i are coefficient matrices and ϵ_t is a zero mean white noise vector with model residuals. This multivariate model depends upon which variables comprises $\mathbf{Y}_t$.

For Atlantic salmon, VAR predictors were chosen based on the main findings of Sandaker et al. (2017) in their recent study of superior Atlantic salmon predictors. The five most significant predictors were: *Lagged values of Atlantic salmon, exchange rates (EUR/NOK), feed consumption, substitute proteins (meat index) and feed prices (fish meal).* All five time series are described in Table (12). Optimal model specification was decided through the AIC for the in-sample period 2006-2018. The optimal lag length p was 4.

For Brent crude, VAR predictors were chosen based on Westgaard et al. (2017)'s superior oil predictors. They concluded that the three most significant predictors were: *The OSX index* [29]*, the credit spread on BB rates high yield bonds and the* 6- *vs.* 1 *mnd NYMEX oil futures spreads* - all of which are financial indicators and reflect supply/demand drivers indirectly. Therefore, we include *OECD inventory volume* and *OPEC production* in order to capture some of these fundamental effects. The AIC suggests that the appropriate lag length p is 1, but we chose p=3 as it provides good fit and increases the dynamic abilities to the benchmark model on longer forecast horizons.

[29] The PHLX Oil Service Index is designed to track the performance of a set of companies involved in the oil service sector.

Table 14: *AIC information criteria: VAR*

IC	Lag length p					
	1	2	3	4	5	6
AIC, Atlantic Salmon	$-2.55e+03$	$-2.56e+03$	$-2.57e+03$	$-2.60e+03$	$-2.59e+03$	$-2.56e+03$
AIC, Brent Crude	$-2.59e+03$	$-2.58e+03$	$-2.58e+03$	$-2.57e+03$	$-2.56e+03$	$-2.54e+03$

D.2 ARIMA(p, d, q)

An ARIMA(p, d, q) (Autoregressive integrated moving average) model is given by

$$(1 - \sum_{k=1}^{p} \alpha_k \mathcal{L}^k)(1 - d)y_t = (1 + \sum_{k=1}^{q} \beta_k \mathcal{L}^k)\epsilon_t \tag{27}$$

where $\alpha_1 \ldots \alpha_p$ and $\beta_1 \ldots \beta_k$ are estimated autoregressive and moving average coefficients respectively. $\mathcal{L}$ denotes the log operator and y_t denotes the dependent variable. The p, d and q denote the number of autoregressive, seasonal and moving average components, respectively. The univariate ARIMA benchmark depends on the degree of persistence in the dependent variable.

D.3 ARIMA$(0, 1, 0)$

The naïve no-change model relates to a special case of an ARIMA model, where the AR coefficient is equal to 1, i.e. a time series with infinitely slow mean reversion. The ARIMA$(0, 1, 0)$ model is then simply given by

$$y_t = y_{t-1} + u_t \tag{28}$$

E State Space Representation

We present the state space representation of the general model presented in Equations (1)-(2):

$$Y_t = \begin{pmatrix} \Lambda & 0 & \cdots & 0 & I & 0 \end{pmatrix} \begin{pmatrix} F_t \\ F_{t-1} \\ \vdots \\ F_{t-p} \\ \epsilon_t \\ \epsilon_{t-1} \end{pmatrix} \tag{29}$$

$$\begin{pmatrix} F_t \\ F_{t-1} \\ F_{t-2} \\ \vdots \\ F_{t-p} \\ \epsilon_t \\ \epsilon_{t-1} \end{pmatrix} = \begin{pmatrix} A_1 & A_2 & \cdots & A_p & 0 & 0 \\ I_r & 0 & \cdots & 0 & 0 & 0 \\ 0 & I_r & & \vdots & \vdots & \vdots \\ \vdots & & \ddots & 0 & \vdots & \vdots \\ 0 & \cdots & 0 & I_r & 0 & \vdots \\ 0 & \cdots & \cdots & 0 & I_n & 0 \\ 0 & \cdots & \cdots & \cdots & 0 & I_n \end{pmatrix} \begin{pmatrix} F_{t-1} \\ F_{t-2} \\ \vdots \\ F_{t-p} \\ \epsilon_t \\ \epsilon_{t-1} \end{pmatrix} + \begin{pmatrix} u_t \\ 0 \\ 0 \\ \vdots \\ 0 \\ e_t \\ 0 \end{pmatrix} \tag{30}$$

Here, $F_t = (f_{1,t}, \ldots, f_{r,t})'$ and $\epsilon_t = (\epsilon_{1,t}, \ldots, \epsilon_{r,t})'$.

For *Application A* we have that:

$$\Lambda = \begin{pmatrix} \Lambda_1^g & \Lambda_1 & 0 & \cdots & 0 \\ \Lambda_2^g & 0 & \Lambda_2 & & \vdots \\ \vdots & \vdots & & \ddots & 0 \\ \Lambda_B^g & 0 & \cdots & 0 & \Lambda_B \end{pmatrix}, \qquad A_q = \begin{pmatrix} a_g^q & 0 & 0 & \cdots & 0 \\ 0 & a_1^q & 0 & & \vdots \\ 0 & 0 & a_2^q & & \vdots \\ \vdots & & & \ddots & 0 \\ 0 & \cdots & \cdots & 0 & a_B^q \end{pmatrix} \tag{31}$$

Here, $\Lambda_b = (\lambda_i, \ldots, \lambda_j)'$ if the variables i through j belong to sector b.

Similarly, for *Application B* we have that:

$$\Lambda = \begin{pmatrix} \lambda_{1,1} & \lambda_{1,2} & \cdots & \lambda_{1,r} \\ \lambda_{2,1} & \lambda_{2,2} & & \vdots \\ \vdots & & \ddots & \lambda_{n-1,r} \\ \lambda_{n,1} & \cdots & \lambda_{n,r-1} & \lambda_{n,r} \end{pmatrix}, \qquad A_q = \begin{pmatrix} a^q_{1,1} & a^q_{1,2} & \cdots & a^q_{1,r} \\ a^q_{2,1} & a^q_{2,2} & & \vdots \\ \vdots & & \ddots & a^q_{r-1,r} \\ a^q_{r,1} & \cdots & a^q_{r,r-1} & a^q_{r,r} \end{pmatrix} \tag{32}$$

We restrict variable i to not to load on factor f by setting $\lambda_{i,f} = 0$.

F The EM Algorithm

As the state variables (x_t) are unobserved, the parameters θ in the state space representation in Equations (29) and (30) are in general not available in closed form, and a direct numerical maximisation of the likelihood is often computationally demanding. Therefore, we estimate these parameters using the Expectation Maximization (EM) algorithm. The EM algorithm was proposed by Dempster et al. (1977) as a general solution to problems for which incomplete or latent data yield a direct maximization of the likelihood function intractable or difficult to deal with otherwise. The basic principle behind the EM algorithm is to write the likelihood as if the data were complete and to iterate between the *Expectation step* where we "fill in" the missing data in the likelihood, namely the unobserved latent factors, and the *Maximization step* where we re-optimize the likelihood. Iterating between these steps, the likelihood is bound to be monotonically improving, and converges towards a local maximum of the likelihood.

We denote the joint log-likelihood of y_t and x_t, where $t = 1, \ldots, T$, by $\mathbb{L}(\theta|Y, X)$, where $Y = [y_1, \ldots, y_T]$ and $X = [x_1, \ldots, x_T]$. Given the available data $\Omega_T \subseteq Y$, the EM algorithm steps can be described in the following way:

- *E-step.* The expectation of log-likelihood conditional on the data is calculated using the model parameter estimates from the previous iteration, θ_k:

$$\mathcal{L}(\theta, \theta_k) = \mathbb{E}_{\theta_k}\left[\mathbb{L}(\theta|Y, X)|\Omega_T\right] \tag{33}$$

  In the case of latent factors, these are estimated in this step.

- *Maximization step.* The model parameters are re-estimated through the maximization of the expected log-likelihood with respect to θ:

$$\theta_{k+1} = \underset{\theta}{\operatorname{argmax}}\, \mathcal{L}(\theta, \theta_k) \tag{34}$$

The estimation problem is now reduced to a sequence of steps, which essentially involves a pass of the *Kalman smoother*

Expectation Step

In the Expectation step (E-step), the EM algorithm finds the expected values of the latent factors x_t which serve as factor estimates. These estimates are found, given the observable data (y_t) and state space model parameters ($\theta_k = (A_k, \Lambda_k, Q_k, R_k)$). In our case, these are found by using the Kalman smoother, which consists of a forward and a backward pass over the observable data, during which it finds the statistically optimal estimates of the latent factors.[30] The forward pass is better known as the Kalman filter.

Given a linear dynamics model

$$y_t = \Lambda x_t + b_t, \qquad b_t \sim N(0, R) \tag{35}$$

$$x_{t+1} = A x_t + w_t, \qquad w_t \sim N(0, Q), \tag{36}$$

that has initial estimates (x_0) of the state variables, and a set of observable values (Y_t), the Kalman filter does a *forward pass* [31] over the data. The state variables have a conditional covariance, P_t, calculated at each time step. For each time step, the Kalman filter does a *measurement update* and a *time update* of the state variables and their covariance matrix.

$$\text{(Measurement update)} \quad x_{t|t} = x_{t|t-1} + K_t(y_t - \Lambda x_{t|t-1}) \tag{37}$$

$$P_{t|t} = P_{t|t-1} - K_t P_{t|t-1} \Lambda^T \tag{38}$$

$$\text{where} \quad K_t = P_{t|t-1} \Lambda^T \left(\Lambda P_{t|t-1} \Lambda^T + R\right)^{-1} \tag{39}$$

$$\text{(Time update)} \quad x_{t+1|t} = A x_{t|t} \tag{40}$$

$$P_{t+1|t} = A P_{t|t} A^T + Q \tag{41}$$

The K_t is often referred to as the *Kalman gain*, and is essentially a weighting that decides how much emphasis should be put on the discrepancy between the observed value (y_t) and the estimated value (Λx_t). In dealing with missing values, i.e. y_t is missing, the Kalman gain is simply set $K_t = 0$. This can be interpreted as the observable value providing no new information to the estimate of x_t. For further derivation and details of the Kalman filter, we refer to Durbin and Koopman

[30]The Kalman smoother is the optimal linear filter in cases where a) the model perfectly matches the real system, b) the entering noise is white (uncorrelated) and c) the covariances of the noise are exactly known.

[31]Processing time-stamped data, starting at time t=1 and going forward to t=T

(2012).

The backward pass of the Kalman smoother provides the optimal estimate of $x_{t|T}$ $(t < T)$ by taking all observations from a fixed interval $y_1, \ldots, y_T$ into consideration. By doing a backward pass[32] over the data, the smoothed estimates $(x_{t|T})$ are calculated:

$$x_{t|T} = x_{t|t} + D_t(x_{t+1|T} - x_{t+1|t}) \tag{42}$$

$$P_{t|T} = P_{t|t} + D_t(P_{t+1|T} - AP_{t+1|t}A^T - Q)D_t^T \tag{43}$$

$$\text{where} \quad D_t = P_{t|t}A^T P_{t+1|t} \tag{44}$$

The D_t is sometimes referred to as the *smoother gain*. A detailed explanation of the Kalman smoother and derivations of the backward recursion equations can be found in Rauch et al. (1965).

Maximization Step

In the maximization step (M-step), the EM algorithm finds the model parameters θ that maximizes the log-likelihood function, given the observable data y_t and the current estimates of the latent factors x_t. That is, the likelihood function $\mathbb{L}(\theta|y_t, x_t)$ is maximized with respect to each of the model parameters, $\theta = (A, \Lambda, Q, R)$. For a full derivation of the log-likelihood function and maximization of it we refer to Bańbura and Modugno (2014). Without any model parameter restrictions and *without* the feature of handling missing values, the model parameter estimates would be:

$$\Lambda_{k+1} = \left(\sum_{t=1}^{T} \mathbb{E}_{\theta_k}\left[y_t {x'}_t | \Omega_T\right]\right)\left(\sum_{t=1}^{T} \mathbb{E}_{\theta_k}\left[x_t {x'}_t | \Omega_T\right]\right)^{-1} \tag{45}$$

$$A_{k+1} = \left(\sum_{t=1}^{T} \mathbb{E}_{\theta_k}\left[x_t {x'}_{t-1} | \Omega_T\right]\right)\left(\sum_{t=1}^{T} \mathbb{E}_{\theta_k}\left[x_{t-1} {x'}_{t-1} | \Omega_T\right]\right)^{-1} \tag{46}$$

$$R_{k+1} = \operatorname{diag}\left(\frac{1}{T}\left(\sum_{t=1}^{T} \mathbb{E}_{\theta_k}\left[y_t {y'}_t | \Omega_T\right] - \Lambda_{k+1}\sum_{t=1}^{T} \mathbb{E}_{\theta_k}\left[x_t {y'}_t | \Omega_T\right]\right)\right) \tag{47}$$

$$Q_{k+1} = \frac{1}{T}\left(\sum_{t=1}^{T} \mathbb{E}_{\theta_k}\left[x_t {x'}_t | \Omega_T\right] - A_{k+1}\sum_{t=1}^{T} \mathbb{E}_{\theta_k}\left[x_{t-1} {x'}_t | \Omega_T\right]\right) \tag{48}$$

[32] Processing time-stamped data, starting at time t=T and going backwards to t=0

The conditional moments of the latent factors (e.g. $\sum_{t=1}^{T} \mathbb{E}_{\theta_k}\left[x_t {x'}_t | \Omega_T\right]$) can be obtained through the Kalman smoother.[33]

When y_t contains missing values, there must be made modifications to the estimates of Λ and R, as these involve the use of y_t. We introduce a diagonal matrix, W_t, of size n with the $i^t h$ diagonal element equal to 0 if $y_{i,t}$ is missing, and equal to 1 otherwise. Including such a selection matrix ensures that only the available data are used in the calculations. As per Bańbura and Modugno (2014), the expressions for Λ and R when y_t includes missing values are:

$$\operatorname{vec}(\Lambda_{k+1}) = \left(\sum_{t=1}^{T} \mathbb{E}_{\theta_k}\left[x_t {x'}_t | \Omega_T\right] \otimes W_t \right)^{-1} \operatorname{vec}\left(\sum_{t=1}^{T} W_t y_t \mathbb{E}_{\theta_k}\left[{x'}_t | \Omega_T\right] \right) \tag{49}$$

$$\begin{aligned} R_{k+1} = \operatorname{diag}\Bigg(\frac{1}{T} \sum_{t=1}^{T} \Big(& W_t y_t {y'}_t {W'}_t - W_t y_t \mathbb{E}_{\theta_k}\left[{x'}_t | \Omega_T\right] {\Lambda'}_{k+1} W_t - W_t \Lambda_{k+1} \mathbb{E}_{\theta_k}\left[x_t | \Omega_T\right] {y'}_t W_t \\ & + W_t \Lambda_{k+1} \mathbb{E}_{\theta_k}\left[x_t {x'}_t | \Omega_T\right] {\Lambda'}_{k+1} W_t + (I - W_t) R_k (I - W_t) \Big) \Bigg) \end{aligned} \tag{50}$$

Further, as we want to impose restrictions on Λ and A we follow the approach of Bork (2009) and Wu et al. (1996). They impose linear restrictions to the state space model of the form $H_\Lambda \operatorname{vec}(\Lambda) = \kappa_\Lambda$ and $G_A \operatorname{vec}(A) = \delta_A$. Combined with the missing values case, the resulting estimates for Λ and A are:

[33] See Watson and Engle (1983)

$$\operatorname{vec}(\Lambda^r_{k+1}) = \operatorname{vec}(\Lambda^u_{k+1}) + \left(\left(\sum_{t=1}^{T} \mathbb{E}_{\theta_k} \left[x_t x'_t | \Omega_T \right] \right)^{-1} \otimes R_k \right) H'_\Lambda$$
$$\times \left(H_\Lambda \left(\left(\sum_{t=1}^{T} \mathbb{E}_{\theta_k} \left[x_t x'_t | \Omega_T \right] \right)^{-1} \otimes R_k \right) H'_\Lambda \right)^{-1} \left(\kappa_\Lambda - H_\Lambda \operatorname{vec}(\Lambda^u_{k+1}) \right) \tag{51}$$

$$\operatorname{vec}(A^r_{k+1}) = \operatorname{vec}(A^u_{k+1}) + \left(\left(\sum_{t=1}^{T} \mathbb{E}_{\theta_k} \left[x_{t-1} x'_{t-1} | \Omega_T \right] \right)^{-1} \otimes Q_k \right) G'_A$$
$$\times \left(G_A \left(\left(\sum_{t=1}^{T} \mathbb{E}_{\theta_k} \left[x_{t-1} x'_{t-1} | \Omega_T \right] \right)^{-1} \otimes Q_k \right) G'_A \right)^{-1} \left(\rho_A - G_A \operatorname{vec}(A^u_{k+1}) \right) \tag{52}$$

Here, A^r_{k+1} and Λ^r_{k+1} are the estimates of the restricted parameters, while A^u_{k+1} and Λ^u_{k+1} refer to the unrestricted parameter estimates from Equations (46) and (49).