Modelling interdependence in a pair of heating oil and natural gas futures curves

Master's Thesis submitted

by

Mikhail Zolotko
(522074)

to

Prof. Dr. Ostap Okhrin

Prof. Dr. Brenda López Cabrera

Ladislaus von Bortkiewicz Chair of Statistics
C.A.S.E.- Centre for Applied Statistics and Economics
School of Business and Economics
Humboldt-Universität zu Berlin

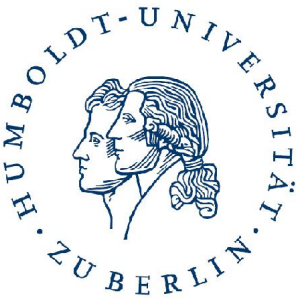

in partial fulfilment of the requirements for the degree of

Master of Economics and Management Science

Berlin, May 29, 2012

Declaration of Authorship

I hereby confirm that I have authored this master thesis independently and without use of others than the indicated resources. All passages, which are literally or in general matter taken out of publications or other resources, are marked as such.

Berlin, 29.05.2012 Mikhail Zolotko

Abstract

This Master thesis further develops the framework for the joint modelling of several futures curves proposed by Ohana (2010). The key innovation is the incorporation of dynamic conditional correlation models based on hierarchical Archimedian copula (HAC-DCC). The conducted analysis allowed to forecast the distribution of the returns of any portfolios composed of the available futures contracts for short time periods. As shown in the study, value-at-risk estimates derived from the forecasts produced by HAC-DCC models are accurate, and these models outperform other benchmark models on a consistent basis as shown by the value-at-risk backtesting procedure carried out on a set of 1000 simulated futures portfolios.

Keywords: multivariate GARCH, hierarchical Archimedean copula, value-at-risk, commodities, forward curves

Contents

List of Figures

List of Tables

1 Introduction

Nowadays more and more agents are involved in trading of energy commodities and various instruments linked to their prices: these agents include banks, electric utilities, commodity producers, specialised trading companies and others. Some of them deal with several instruments that may based only on one commodity, but have quite different properties. Nevertheless, the situation becomes even more complicated if an agent has to deal with instruments related to several commodities whose prices may be partly influenced by common factors. Risk management and portfolio optimisation in such multi-product and especially multi-commodity setting where product prices may be dependent on each other in a complicated way is a challenging task and requires reliable tools for pricing various instruments.

Commodity futures and forward contracts play a special role in the world of the energy-related instruments. First, they represent the simplest type of commodity derivatives, an agreement to buy or sell the underlying asset at a predetermined time in the future for a particular price that is agreed today. Their difference is that the former is a standardised contract that is traded on a futures exchange, whereas the latter is an OTC product that does not have to be standardised, but rather can take into account specific needs of the buyer and/or the seller. However, as noted in Pilipovic (2007), due to the nature of the energy commodity market, the terms forward and futures price can be used interchangeably, if the delivery and the payment dates of both contracts coincide and "there is no possibility of default on either side" because they represent the same value. The second feature of the forward and futures contracts is that despite their simplicity, the so-called forward curves formed by all available futures/forward prices at a particular point in time, constitute an essential input to the pricing model of any other energy derivatives. Hence the importance of the ability to model the whole set of futures/forward prices stems from the two sources: 1) due to their simplicity, forward and futures contracts are among the most popular and liquid energy-related instruments which means that many market participants have open positions in the futures/forward markets and are interested in tools allowing to predict the short- (or possibly long-) term dynamics of their futures/forward portfolios, 2) a reliable model of the futures price dynamics allows to price other more complex commodity derivatives.

In this thesis, the first motivation is prevailing. We assume that modelling the dynamics of a commodity curve is a self-sufficient goal, but the techniques and methods

applied in the work may also find their application in the second line of possible research, i.e. pricing of more complex derivatives.

Throughout the paper we will use the notation $F(t,T)$ for the observed at time t price of a forward/futures contract with maturity (which we assume to coincide with the last trading day) at time T, $T > t$. Let T be the time of maturity of some futures contract in general, additionally let T_t^i denote the last trading day of the i-th nearby futures observed on day t. There seems to be a convention to refer to a commodity price curve as a forward curve, i.e. not a futures curve. This may be reasonable since forward and futures price typically refer to the same value. However, in this thesis we decided to refer to the forward curve as a futures curve to avoid any confusion and to emphasise that market futures prices were used in the empirical work. The term forward curve will be encountered only in the review of the previous research.

As mentioned above, a multi-commodity setting presents an additional challenge for practitioners. Ohana (2010) is a relatively recent study that proposed a framework of evolution of the heating oil and natural gas futures curve. An attractive feature of this study is a combination of elements of several approaches to the commodity spot and futures price modelling (two-factor model of the futures curve dynamics, vector error-correction model, copula-based multivariate distributions). However, the solid framework was applied to draw only descriptive conclusions and was not used to solve any problems that may be faced e.g. by a commodity risk manager.

This thesis was inspired by Ohana (2010) and follows the basic structure of this paper. In the first step, we calibrate the so-called two-factor model that helps describe the whole range of futures returns of the two commodities with only four series of shocks. Next, in line with Ohana (2010), the deterministic component is extracted from the shocks series by estimating a model of the vector autoregression class. After that, the series of the residuals in this model are analysed. Our main contribution is the estimation of the multivariate GARCH models based on the hierarchical Archimedian copula, a recent development in the GARCH family which is a flexible instrument allowing for a variety of possible dependence structures for several time series. The ultimate goal of estimating these models is to predict value-at-risk for a particular futures portfolio for a given day, which is derived from the predicted distribution of the return of this portfolio for this day. The models are estimated on many time windows, and value-at-risk is estimated for 1000 portfolios. Value-at-risk backtesting and subsequent benchmarking with other nested and non-nested model specifications showed that the incorporation of multivariate GARCH models based on the hierarchical Archimedian copula can make the framework of Ohana (2010) a useful risk management tool.

This thesis is organised as follows. Chapter 2 reviews approaches to modelling future prices developed to date, classifies them into two groups and shows the theoretical

link between them. Chapter 3 focuses on the two-factor model of the futures curve described in Ohana (2010). Chapter 4 provides necessary theoretical background on multivariate GARCH models including those based on hierarchical Archimedian copulas. Chapter 5 is a simulation study where we investigate how competing model specifications can approximate data generated by each other and what the consequences of a model misspecification are. Chapter 6 presents the multi-stage empirical study. Chapter 7 concludes.

2 Review of approaches to futures price modelling

All approaches to model futures curves can be subdivided into two broad groups: 1) theoretical stochastic models for a spot price that provide a framework for pricing various commodity derivatives including futures and 2) methods to model and analyse futures curves as a whole (see e.g. Clewlow and Strickland (2000)).

2.1 Stochastic models for spot prices

Development of stochastic methods for the modelling of commodity prices began in the late 1980s – early 1990s and was inspired by the yield curve modelling techniques developed by that time. The main assumption in such models is the process followed by the spot price of a commodity. Clewlow and Strickland (2000) and Geman (2005) mention several models of this kind having different number of state variables. Simplest models assume only one state variable which is the current spot price itself. The general form of the diffusion process followed by the spot price in this case is:

$$dS(t) = \mu\{S(t)\}\, dt + \sigma\{S(t)\}\, dW(t), \tag{2.1}$$

where $S(t)$ is the spot price, t is the time of observation, $\mu\{S(t)\}$ is the drift rate, $\sigma\{S(t)\}$ is the instantaneous volatility of the spot price, $W(t)$ is the Wiener process. Model 1 in Schwartz (1997) is an example of this general formulation. In this model, the drift part $\mu(S_t)dt$ exhibits mean reversion to some long-term value, and the variance is proportional to the current spot price. The risk-neutral dynamics of the spot price has the form:

$$dS(t) = h\{\theta - \nu - \ln S(t)\}\, S(t)dt + \sigma S(t)dW(t), \tag{2.2}$$

where θ is the long-term value of the spot price, h is the speed of mean-reversion, ν is the market price of energy risk and σ is the constant volatility of the spot price return. Schwartz (1997) shows that under such assumptions the forward price will be expressed as:

$$F\{t, T, S(t)\} = A(t,T)\{S(t)\}^{B(t,T)}, \tag{2.3}$$

where

$$A(t,T) = \exp\left[\left\{1 - e^{-h(T-t)}\right\}\alpha + \frac{\sigma^2}{4h}\left\{1 - e^{-2h(T-t)}\right\}\right],$$
$$B(t,T) = e^{-h(T-t)},$$
$$\alpha = \theta - \frac{\sigma^2}{2h} - \nu.$$

This simple model can be extended by introducing seasonality (although this is not relevant for all commodities) as in Borovkova and Geman (2008) or by the inclusion of jumps (see Clewlow and Strickland (2000)) as previously done by Merton (1976) for the purposes of stock option pricing.

More sophisticated models introduce more state variables in addition to the spot price, all of which are assumed to follow a joint stochastic process. Apart from the simple mean-reversion model discussed above, Schwartz (1997) also discusses two its generalisations: the two-factor model with spot price and net convenience yield being the factors, and the three-factor model that additionally includes instantaneous interest rate as the third factor. The two factors in the models described in Gabillon (1991) and Pilipovic (2007) are spot price and long-term value of mean-reversion. Cortazar and Schwartz (2003) propose a model with three factors: spot price, net convenience yield and long-term spot price return. The factors in the model of Eydeland and Geman (1998) are spot price and instantaneous volatility, analogous to the interest rate model of Heston (1993). One of the examples of recent developments in this area of research is Liu and Tang (2010) who criticise the idea of modelling net convenience yield (convenience yield minus storage costs), i.e. implicit treating of storage costs as a percentage of the spot price and suggest treating convenience yield and storage costs separately instead.

In most cases, estimation of the parameters in these models implies the using the Kalman filter since the state variables are not directly observable (see an example of the Kalman filter implementation in Schwartz (1997)). In some cases it is possible to use procedures that do not impose strict requirements on the data with regard to the missing values and are considerably easier to implement as compared to the Kalman filter. Thus, Cortazar and Schwartz (2003) applied an iterative procedure to their three-factor model and to the two-factor model of Schwartz (1997) to estimate both the factors and the values of the state variables for each observation date. These estimates were shown to be "reasonably close" to those obtained using the Kalman filter.

One of the central questions in stochastic modelling of forward or futures commodity prices is formulated in Gabillon (1991): he pointed out an important difference between the goal of *"developing a model which describes the motion of the term structures of futures prices and volatilities with satisfactory accuracy"* and *"developing a model that*

adequately values most of the derivative securities". Both goals are important to a risk manager. But in fact, there exists a trade-off between them - allowing some parameters to be time-varying one could develop a model that would price most traded derivatives on a particular commodity fairly correctly. However, such a model would normally be just a static fit with poor dynamic properties. In order for the model to have adequate dynamic behaviour, one should either assume all its parameters to be constant or to be a very simple function of time.

2.2 Forward curve models

Specifying spot commodity price dynamics is not an essential step for the modelling of commodity forward/futures prices. It is also possible to model the dynamics of the latter explicitly, defining risk factors only in general form. This approach also allows to model all forward/futures prices simultaneously. One of the first studies to formulate the framework for this approach was Reisman (1991). The risk-neutral process followed by the futures price $F(t,T)$ is assumed to have the form:

$$\frac{dF(t,T)}{F(t,T)} = \sum_{l=1}^{d} \chi_l(t,T) dW_l(t), \tag{2.4}$$

where $W_1, W_2, ..., W_d$ are d independent Brownian motions and $\chi_l(t,T)$ are the corresponding volatility functions of the futures prices, or risk factor loadings.

Under the risk-neutral probability measure, a position in an asset that requires no initial investment, such as a futures contract, must have a zero expected return. Therefore the dynamics of the futures prices in (2.4) has zero drift.

As described in Clewlow and Strickland (2000), the next step in this framework is to define the number of factors and to estimate the corresponding volatility functions $\chi_l(t,T)$. In general case, $\chi_l(t,T)$ are not assumed to have any parametric form. To define the factors, many authors follow the methodology of Heath et al. (1990) and apply principal components analysis (PCA) which implies the eigenvalue decomposition of the sample covariance matrix $\widehat{\Sigma}$ of all the futures prices $F(t,T)$ for all M available maturities, whose elements $\widehat{\sigma}_{ij}$ are estimated as:

$$(\widehat{\sigma})_{ij} = \frac{1}{N} \sum_{t=1}^{N} (\widehat{r}_t^i - \overline{\widehat{r}}^i)(\widehat{r}_t^j - \overline{\widehat{r}}^j), \tag{2.5}$$

where $\widehat{r}_t^i$ is the observed at time t return of the i-th nearby futures, $\overline{\widehat{r}}^i$ is a sample mean of the return of the i-th nearby futures, N is the number of observations for each

considered maturity, $i, j = 1, \ldots, M$. The matrix $\widehat{\Sigma}$ is decomposed as:

$$\widehat{\Sigma} = \Gamma \Lambda \Gamma^{\top}, \tag{2.6}$$

where the columns of Γ are the eigenvectors of $\widehat{\Sigma}$ and Λ is the diagonal matrix composed of the corresponding eigenvalues of $\widehat{\Sigma}$:

$$\Gamma = \begin{pmatrix} \gamma_{11} & \gamma_{12} & \cdots & \gamma_{1M} \\ \gamma_{21} & \gamma_{22} & \cdots & \gamma_{2M} \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_{M1} & \gamma_{M2} & \cdots & \gamma_{MM} \end{pmatrix} \text{ and } \Lambda = \begin{pmatrix} \lambda_1 & 0 & \ldots & 0 \\ 0 & \lambda_2 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & \lambda_M \end{pmatrix}.$$

The share of the variance υ explained by the first d principal components in the total variance can be calculated as:

$$\upsilon = \frac{\sum_{l=1}^{d} \lambda_l}{\sum_{i=1}^{M} \lambda_i}. \tag{2.7}$$

In practice, it is possible to explain a considerably high share of variance with d first factors where $d < M$ (see below), which allows to reduce the dimensionality of the subsequent analysis by disregarding factors with low explanatory power. The estimates of the volatility functions in (2.4) are obtained as $\chi_l(t, T_t^i) = \gamma_{il}\sqrt{\lambda_l}$.

Clewlow and Strickland (2000) showed that similar to the derivation of the forward price from a spot price diffusion process such as given in (2.1), one can also derive the process for the spot price from the futures price process, such as provided in (2.4). Moreover, they showed that the one-factor version of 2.4 with volatility of the factor parametrically defined as $\chi_1(t, T) = \chi e^{\kappa(T-t)}$ with χ and κ being parameters, is equivalent to the single factor Schwartz (1997) model (2.2) with time dependent term $(\theta - \nu)$.

Cortazar and Schwartz (1994) is one of the first studies where PCA is applied to the commodity futures curves analysis in order to define the optimal number of factors and to estimate the corresponding volatility functions in (2.4). Using historical data, the authors simulated stochastic processes of the copper futures prices which allowed them to estimate expected cash flows of a copper-linked note and define its price.

Cortazar and Schwartz (1994) is also one of the first works where the first three determined factors defining the dynamics of a commodity forward curve were denoted as "level", "steepness" and "curvature". These terms were used in a somewhat abstract sense since e.g. by "level" no parallel shift is meant and "steepness" does not correspond to any commonly used steepness measure. Nevertheless, these factors turned out to be more convenient for modelling yield curve dynamics than more "intuitive" ones. Cortazar and Schwartz (1994) note that the patterns of the factor loadings of the principal components resulted from their analysis resemble those of the factors defining

the dynamics of a yield curve described in the earlier study Litterman and Scheinkman (1991).

A more recent study Tolmasky and Hindanov (2002) attempts to consider heating and crude oil futures traded on NYMEX within one model, i.e. to define factors affecting the shape of both curves. PCA in this case is performed on the whole dataset, and the seasonality problem which is present in the case of the heating oil is taken into account. Factor loadings are corrected for each season depending on the variance explained by a particular factor in the past seasons. However, the authors admit that it is not clear how to assess the statistical significance of the seasonality in this case. The estimated model is used to simulate futures contract returns which allows to estimate value-at-risk and price complex commodity derivatives such as basket spread options.

Notably, the explanatory power of the factors governing the dynamics of the futures curve of a single commodity is reported to be remarkably high - even two factors explain around 96-97% of the variance as documented by Cortazar and Schwartz (1994). When applied to the explanation of the dynamics of more than one commodity curve within one model, as in Tolmasky and Hindanov (2002), PCA yields worse results in terms of the variance explained by the factors. Besides, the factors become less interpretable. These results are confirmed in Chantziara and Skiadopoulos (2008) who consider futures on four different oil products. They also note that the quality of the regression models where futures prices are explained by the revealed factors does not improve after applying PCA to all contracts simultaneously.

Koekebakker and Ollmar (2005) who studied the dynamics of electricity prices on the Nord Pool electricity derivative exchange documented much lower explanatory power of the principal components: 75% of the variance is explained by two factors whereas more than 10 factors are required to explain 95% of the variance. This may be attributed to the higher complexity of the electricity price dynamics and a specific nature of this market.

Järvinen (2004) studied the dynamics of prices for Brent oil and NBSK using PCA and found a relatively poor explanatory power of the principal components (two factors explain around 63% of the variance for NBSK and 81% for Brent, four factors are required to explain more than 90% of the variance for both commodities). He also notes that the components themselves cannot be cleary interpreted since they do not look similar to the typical "level", "steepness" and "curvature" factors. It is necessary to point out that in contrast to many other studies, he used broker's swap quotes as the source of the price data which may have influenced the obtained results.

Model (2.4) can also be formulated under the physical probability measure $\mathbb{P}$. But in this case, since a futures contract is a risky investment, it does not necessarily have a zero expected return which means that the terms $W_l(t)$ are not necessarily Wiener processes

under $\mathbb{P}$. For the purposes of the empirical work the model is usually discretised, so the times series formed by the terms $dW_l(t)$ can be broadly denoted as shocks series and modelled separately.

An example of forward curve modelling under $\mathbb{P}$ is the study Borak and Weron (2008). Prompted by the unsatisfactory results yielded by PCA for electricity markets they applied more sophisticated dynamic semiparametric factor model (DSFM) that also fits the general framework of (2.4). Analysing the Nord Pool dataset also used in the study Koekebakker and Ollmar (2005) mentioned above, Borak and Weron (2008) found that DSFM is "an efficient tool for approximating forward curve dynamics" in application to the electricity markets.

The two-factor model used in the recent work Ohana (2010) also defines the commodity process under the physical probability measure and employs parametrically defined factor loadings. This model is used as a first step to get insights into the dynamic interdependence of the heating oil and natural gas futures curves in the US market. As the motivation to investigate the interrelations between the two markets, economic reasons, such as links on both demand and supply side of the markets were cited. Thus, on the demand side there is a certain switching potential between the two commodities in the industry sector, power generation sector and at times even in the private sector. Since the two commodities are substitutes to some extent, if the price for one of the commodities is driven up by external factors, the agents switching to the cheaper commodity will put upward pressure on its price, too. On the supply side the link between the markets is not so well defined, i.e. there exist at least two effects that may influence the price relationship in the opposite directions. As the reason forcing the prices to co-move, the geographical position of major US gas and oil fields (the Gulf of Mexico) is mentioned. This area is often adversely affected by weather conditions and even natural catastrophes that cause a rise in both commodity prices.

The factors defining the shape of the futures curve every day, "level" and "steepness" (called "slope" in the paper), are predetermined in the model framework. Ohana (2010) uses the two-factor framework to transform the dataset of 28 futures return series to the 4-variate set of so-called long-term and short-term price shocks of the two commodities. The four series of shocks are used in the subsequent time series analysis to discover both long- and short-term dependence structures.

The idea to decompose "daily deformations" of a futures curve into short-term and long-term shocks had already been explored in e.g. Manoliu and Tompaidis (2002), Schwartz and Smith (2000) and Geman and Nguyen (2005). Besides, in connection with modelling of the yield curve dynamics, these concepts have an even longer history and were mentioned already in Heath et al. (1990). However, the advantage of the model used in Ohana (2010) is the combination of straightforward interpretability with

a good explanatory ability: the percentage of explained futures return variance is over 95%. A detailed description of the model is given in chapter 3.

Ohana (2010) was not the first to study interrelations of the futures prices for different commodities. The innovation of this study is the combination of an analytical approach to the futures curve decomposition with the subsequent cointegration and error-correction time series analysis in the spirit of e.g. Asche et al. (2006), Bachmeier and Griffin (2006), Dawson et al. (2006), Grasso and Manera (2007) or Hartley et al. (2008). The mentioned authors pay a lot of attention to the econometric model specification in order to find out laws of market interactions. But they use only front-month futures prices which probably constrains the power of the techniques applied, whereas the work of Ohana (2010) makes use of the wide spectrum of the available futures prices.

Ohana's (2010) study resolves among other things a problem that some researchers seem to ignore: on most futures markets one can never observe one and the same futures curve on different days because the spectrum of available maturities changes every day (i.e. all maturities decrease by one day). This may pose a problem if a futures contract has a long (say, monthly) delivery cycle. In this case, nearby futures prices shortly after the delivery (with maturity in a little less than one month) will be regarded exactly in the same manner as nearby futures prices shortly before the delivery. Ohana's model automatically accounts for these differences in remaining days to maturity.

As mentioned in chapter 1, in our research we follow the general structure of Ohana (2010), but use recent developments in the multivariate GARCH modelling to better capture the variance-covariance of the series vector in question. Besides, we pursue a practical goal to find a reliable way to produce accurate value-at-risk forecasts for various portfolios composed of futures contracts.

3 Two-factor model of the futures curve (Ohana, 2010)

This chapter is devoted to the description of Ohana's (2010) two-factor model. Its main features are: 1) it explicitly defines the factors that determine the shape of the futures curve at every moment t and 2) it enables decomposition of daily futures curve moves into long-term shocks that affect all maturities equally, and short-term shocks that affect shorter maturities more significantly that longer ones. As economic justification for this decomposition the following interpretation of shocks is given in Ohana (2010):

- the short-term shocks are related to events that are expected to affect the market for a limited period of time, such as temperature change, transitory supply shortage, transportation problems etc.
- the long-term shocks reflect fundamental changes in supply or demand for the commodity, such as new information on available reserves, change of political situation in commodity-rich countries etc.

The model is defined in discrete time on a filtered probability space $(\Omega, \mathcal{F}_t, \mathbb{P})$ where the probability measure $\mathbb{P}$ is physical. For $t \geq 0$ the following arbitrage-free dynamics for the futures prices of a commodity is assumed:

$$\frac{F(t,T) - F(t-\Delta,T)}{F(t-\Delta,T)} = \exp\left\{-\frac{k}{252}(T-t+\Delta)\right\}\delta_t^S + \delta_t^L, \tag{3.1}$$

with Δ being some small time interval, $(\delta^S)_{t\geq 1}$ and $(\delta^L)_{t\geq 1}$ are $\mathcal{F}_t$-adapted short-term and long-term shocks which we decompose as

$$\begin{aligned} \delta_t^S &= \lambda_t^S + \xi_t^S, \\ \delta_t^L &= \lambda_t^L + \xi_t^L, \end{aligned} \tag{3.2}$$

where λ_t and ξ_t are the deterministic and random shocks components respectively, k is the characteristic value of a commodity defining the extent to which a short-term shock affects longer maturities and 252 is the assumed number of trading days in a year.

In (3.1) and thereafter, T, t and Δ are measured in trading days. In the subsequent analysis Δ is equal to one trading day.

Both deterministic and random components of the shocks are in general mutually dependent. The general model view defined by (3.1) is assumed to remain unchanged throughout the whole time period considered in the study for both commodities. In particular, the parameter k is assumed to be constant for each commodity. Contrary to this, the mechanisms defining the behaviour of the shocks components are not required to be stable over the course of time. Econometric models that parametrise various types of the dependence of the shocks components constitute the essence of this study. After both futures curves datasets are transformed into the shocks series, the next step is to extract the deterministic component using the VAR framework, and then apply different specifications of multivariate GARCH to capture the residual variance.

Every day t we observe a futures curve which consists of the prices of the first nearby (soonest to mature), second nearby, ..., M-th nearby futures contracts. Hence, M futures returns can be calculated every day. In order to decompose the futures price returns into shocks consistently, these returns need to be consistently calculated, i.e. only the price of one and same contract can be used to calculate a return. For example, if today is the last day when the first nearby futures contract is traded, tomorrow's return of the first nearby futures has to be calculated using the tomorrow's and today's price of the futures contract which is still second nearby today but will the be the first nearby tomorrow. Let us denote $\widehat{r}_t^i$ the return of the i-th nearby futures calculated on day t.

In order to be able to transform futures returns into shocks, the model must be calibrated. A crucial assumption that considerably facilitates the calibration of the model is that the short-term shock does not affect the farthest (M-th) return which can be formally written as $\exp\left\{-k(T_t^M - t + \Delta)/252\right\} \approx 0$. The most distant futures return is therefore fully attributed to the long-term shock which implies that the short term shock can be easily expressed e.g. from the actual return of the first nearby futures. The expressions for both shocks are thus:

$$\begin{aligned} \delta_t^L &= \widehat{r}_t^M, \\ \delta_t^S &= \exp\left\{\frac{k}{252}\left(T_t^1 - t + \Delta\right)\right\}\left(\widehat{r}_t^1 - \widehat{r}_t^M\right). \end{aligned} \tag{3.3}$$

Plugging these expressions for shocks to (3.1), we obtain the following formulation for the model-implied return (denoted by r_t^i) of the futures contract with the maturity T_t^i that can be written as a function of k:

$$r_t^i(k) = \widehat{r}_t^M + \exp\left\{-\frac{k}{252}\left(T_t^i - T_t^1\right)\right\}\left(\widehat{r}_t^1 - \widehat{r}_t^M\right). \tag{3.4}$$

In order to calibrate the model, i.e. estimate k, we minimise the root mean square error (RMSE) calculated from squared differences between the model-implied (r_t^i) and actual

$(\widehat{r}_t^i)$ futures returns using all available information (every trading day across all available maturities) for each commodity:

$$\widehat{k} = \arg\min_k RMSE = \sqrt{\frac{1}{NM}\sum_{t=1}^{N}\sum_{i=1}^{M}\left\{\widehat{r}_t^i - r_t^i(k)\right\}^2}, \tag{3.5}$$

where N is the number of returns observations (which requires $N+1$ futures curve observations) and M is the number of considered maturities.

An attractive and desirable for us feature of model 3.1 is that it can also be interpreted as a model for the shape of a futures curve after some rearranging of terms. This also implies that discrete daily returns are actually an approximation of continuous returns:

$$\Delta \log F(t,T) \approx \frac{F(t,T) - F(t-\Delta,T)}{F(t-\Delta,T)} = \exp\left\{-\frac{k}{252}(T-t+\Delta)\right\}\delta_t^S + \delta_t^L, \tag{3.6}$$

which repeats (3.1) and equates it to the difference of the logarithmised futures prices. Isolating $\log F(t,T)$, we obtain the expression for the futures curve at time t:

$$\log F(t,T) \approx \log F(0,T) + \sum_{j=1}^{t}\exp\left\{-\frac{k}{252}(T-j+\Delta)\right\}\delta_j^S + \sum_{j=1}^{t}\delta_j^L; \quad t \in [0;T]\,. \tag{3.7}$$

Now, if we assume the initial futures curve to take the form:

$$\log F(0,T) = Q(T) + \exp\left\{-\frac{k}{252}T\right\}X_0 + Y_0; \quad T \geq 0, \tag{3.8}$$

where X_0 and Y_0 are real deterministic numbers and $Q(T)$ is some periodic function with the 12 months cycle that ensures the seasonality property of a futures curve observed in chapter 6, then (3.7) can be rearranged as:

$$\log F(t,T) \approx Q(T) + \exp\left\{-\frac{k}{252}(T-t)\right\}X_t + Y_t; \quad t \in [0;T]\,, \tag{3.9}$$

where for $t \geq 1$, X_t and Y_t are defined respectively as:

$$X_t = X_0\exp\left(-\frac{k}{252}t\right) + \sum_{j=1}^{t}\exp\left\{-\frac{k}{252}(t-j+\Delta)\right\}\delta_j^S, \tag{3.10}$$

$$Y_t = Y_0 + \sum_{j=1}^{t}\delta_j^L. \tag{3.11}$$

Equation 3.9 shows that a futures curve is decomposed into two summands: Y_t which regulates the vertical shift of the curve (later referred to as "level") and the product of

two factors, one of which, $\exp\{-k(T-t)/252\}$, depends only on the time to maturity for a particular commodity curve on a given day. Hence the exact shape of the curve is governed solely by X_t which we will refer to as "slope". It is easy to see that positive values of X_t will generate a curve that declines with increasing time to maturity (the situation referred to as backwardation) and, conversely, negative values of X_t generate an upward-sloping curve (which is referred to as contango).

At this point it is worthwhile to note that although the described model can generate futures curves that exhibit the seasonality property, the presence of this property itself is irrelevant to the subsequent analysis. Since the value of the seasonality function $Q(T)$ is constant for every futures price for some T, its value cancels out when futures price returns are calculated which is why no seasonality term shows up in (3.1) and all other expressions that directly follow from it.

It is clear that the names of the factors defining the shape of a futures curve in the described model (slope and level) are somewhat allegoric. Nevertheless, to a considerable extent these factors are analogous to purely empirical principal components such as those documented by Cortazar and Schwartz (1994). It is even possible to claim that "slope" in the Ohana's model corresponds to both "steepness" and "curvature" regarded from the common sense point of view. However, one must understand that one factor cannot be so flexible and cannot allow for arbitrary separate changes in "curvature" and "steepness", e.g. the Ohana's model cannot produce a straight line (a line with "zero curvature") other than the one parallel to the horizontal axis.

Ohana (2010) also provides fundamental theoretical interpretation of the factors in the light of the risk factors in spot commodity price models reviewed in chapter 2. "Level" corresponds to the long-term commodity price and "slope" is related to the convenience yield, i.e. to the relative benefit of holding of a physical commodity as compared to having a long futures position in this commodity.

Finally, we note that the slopes and the levels of the futures curves can be easily estimated from the futures price data. Assuming the level to be zero on the day of the first observation, we can use (3.11) to derive the levels for all other days. The slopes can be estimated applying (3.9) for some maturities with the same value of the seasonal function $Q(T)$ (e.g. the 1st and the 13th) and subtracting one expression from the other. Assuming that $\exp\{k(T_t^{13}-1)/252\} \approx 0$, we obtain the following approximation for the slope on day t:

$$X_t \approx \exp\left\{\frac{k}{252}(T_t^1 - t)\right\} \log\left(\frac{F(t,T_t^1)}{F(t,T_t^{13})}\right). \tag{3.12}$$

The shocks will represent our time series of interest in the following analysis described in chapter 6. Since we are dealing with two commodities, the subsequent analysis implies four-dimensional models.

4 Copula-based multivariate GARCH

Copula-based multivariate generalised autoregressive conditional heteroskedasticity model (C-MGARCH) was proposed by Lee and Long (2009) as a generalisation of the family of multivariate generalised autoregressive conditional models (MGARCH), such as the BEKK model (Engle and Kroner (1995)), the DCC model (Engle (2002)) or the VC model (Tse and Tsui (2002)). The main advantage of such generalisation is a possibility to loosen assumptions regarding the distribution of the error terms. More specifically, the error terms are allowed to be conditionally non-normally distributed which is line with the empirical studies Fama and French (1993), Richardson and Smith (1993), Longin and Solnik (2001) and many others that reject the conditional multivariate normality assumption. In C-MGARCH, the error terms are uncorrelated, but can still be (nonlinearly) dependent which represents a situation that is impossible under the multivariate normality assumption for the error distribution. The nonlinear dependence structure of the error terms is controlled by a copula. Thus, another side of the looser requirements imposed on error terms in C-MGARCH is the possibility to separately model both nonlinear dependence of the time series through various copula specifications and their linear dependence by means of the standard MGARCH.

In this study we employ hierarchical Archimedian copulas (HAC) that provide an especially flexible way to capture different dependence structures using a low number of parameters which proves very useful in the high-dimensional setting. Later we compare the performance of MGARCH models based on the hierarchical Archimedian copulas (HAC-MGARCH hereafter) with that of those based on plain Archimedian copulas (AC-MGARCH hereafter) and those employing no copula (standard MGARCH hereafter). In this chapter, first essential theoretical background on copulas is provided and the concept of HAC is explained, then standard MGARCH framework is presented. Finally, it is shown how the two approaches are combined in C-MGARCH models.

This chapter presents general C-MGARCH framework and describes an approach to the estimation of all models of this family. The exact lists of models used in the simulation and empirical studies are provided in the respective chapters.

4.1 Copulas

The concept of copula was introduced in the seminal paper of Sklar (1959). An extensive introduction into the properties of copulas is provided in Nelson (2006) and Joe (1997). By definition, a copula is a function that helps describe some multivariate distribution by formalising a relation between the values of the marginal distributions and the values of the corresponding multivariate distribution. Formally, a copula is a continuous function $C\colon [0;1]^d \to [0;1]$ that satisfies the equality:

$$F(x_1, \ldots, x_d) = C\{F_1(x_1), \ldots, F_d(x_d)\}, \; x_1, \ldots, x_d \in \mathbb{R}, \tag{4.1}$$

where $F(x_1, \ldots, x_d)$ is a d-dimensional multivariate distribution, and $F_1(x_1), \ldots, F_d(x_d)$ are the respective marginal distributions.

In his theorem, Sklar (1959) showed the function C to be unique for every continuous multivariate distribution. By its nature, every copula itself is also a multivariate distribution function with uniform margins. In practice, a class of so-called Archimedian copulas proved to be very convenient due to its useful properties: Archimedian copulas can generate a big variety of dependence structures, and a closed form is available for all copulas of this class. In general case a d-dimensional Archimedian copula has the form:

$$C(u_1, \ldots, u_d) = \phi^{-1}\{\phi(u_1) + \ldots + \phi(u_d)\}, \; [u_1, \ldots, u_d] \in [0, 1], \tag{4.2}$$

where ϕ is a copula generator function. Both ϕ and ϕ^{-1} need to have a closed form for (4.2) to have a closed form as well. Necessary and sufficient properties of ϕ are provided in McNeil and Nešlehová (2009). Some of the most widely used generator functions are the Gumbel generator $\phi_\theta(u) = (-\log x)^\theta$ for $x \in [0, \infty)$, $\theta \in [1, \infty)$ that generates a copula with upper tail dependence:

$$C^{Gumbel}(u_1, \ldots, u_d; \theta) = \exp\left[-\left\{(-\log u_1)^\theta + \ldots + (-\log u_d)^\theta\right\}^{1/\theta}\right], \tag{4.3}$$

and the Clayton generator $\phi_\theta(u) = \left(x^\theta - 1\right)/\theta$ for $x \in [0, \infty)$, $\theta \in (-1/(d-1), \infty)$, $\theta \neq 0$ that generates a copula with lower tail dependence:

$$C^{Clayton}(u_1, \ldots, u_d; \theta) = \left(u_1^\theta + \ldots + u_d^\theta - 1\right)^{-1/\theta}. \tag{4.4}$$

The concept of copula also nests the case of independence of the variables which is described by the product copula $C(u_1, \ldots, u_d) = u_1 \cdot \ldots \cdot u_d$.

Archimedian copulas are permutation symmetric which can be too restrictive even in the 3-dimensional case. A generalisation of an Archimedian copula called hierarchical Archimedian copula (HAC) solves this problem. A hierarchical copula is a copula that

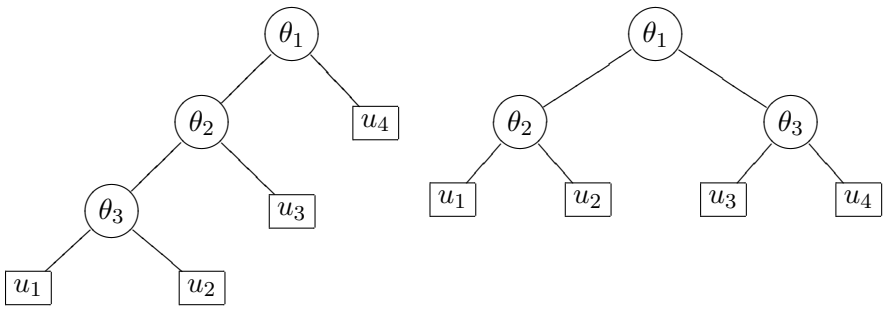

Figure 4.1: Graphical representation of two HAC structures for HAC dimension $d = 4$: s=(((12)3)4) (left) and s=((12)(34)) (right)

may nest other copulas using them as arguments, whereas some of these nested copulas (subcopulas) may in turn nest other copulas, too, thus forming several levels of hierarchy. Hereafter we will refer to the previously described one-parameter copula as plain copula. An extensive overview of properties of HAC is given in Okhrin et al. (2012). The variety of distributions that can be described by a HAC stems from the following sources: copula's structure, employed generator functions and dependence strength reflected in parameter values. There exist certain limitations as to which generator functions (i.e. which copula types) can be combined, and certain conditions that should be imposed on parameters to ensure that the resulting HAC is indeed a copula. These conditions are investigated in McNeil (2008) and Hofert (2008). In this study we will use only the HAC with the same kind of the generator function in all subcopulas. For this case the necessary condition for a function to be a copula is that any outer copula, i.e. a copula nesting another copula, should have a parameter not higher that any of its subordinated (inner) copulas. In other words, the strength of the relationship described by the outer copula should exceed none of those described by any of its subordinated copulas. For the sake of brevity, we will say that copula C_a links variables x_j and x_k and subcopula C_b to mean that $F_j(x_j)$, $F_k(x_k)$ and subcopula C_b are the arguments of the copula C_a.

The structure s of a HAC will be described as $s = \{(\ldots(\ldots i_j \ldots i_k \ldots)\ldots(\ldots)\ldots)\}$ or $s = \left\{\left(\ldots\left(\ldots \text{“}\eta_{i_j}\text{”} \ldots \text{“}\eta_{i_k}\text{”} \ldots\right)\ldots(\ldots)\ldots\right)\right\}$, where “$\eta_{i_\ell}$” are names of the variables following the d-dimensional distribution whose dependence structure is defined by the HAC, $i_\ell \in \{1, \ldots, d\}$ are the indices of these variables, and every expression in brackets describes the structure of a subcopula by showing what variables and/or other subcopulas it links with each other. Figure 4.1 illustrates this notation demonstrating two possible HAC structures.

4.2 Multivariate GARCH

If ξ_t is a d-variate vector with $\mathsf{E}(\xi_t|\mathcal{F}_{t-1}) = 0$, where $\mathcal{F}_{t-1}$ is the information set available at $t-1$, MGARCH models describe the conditional dynamics of its variance-covariance matrix $H_t \stackrel{\text{def}}{=} \mathsf{E}(\xi_t\xi_t^\top|\mathcal{F}_{t-1})$. It is natural to measure time in trading days as in chapter 3 since daily data will be analysed in the empirical part of the study.

Vector ξ_t is expressed as $\xi_t = H_t^{1/2}e_t$, where $e_t \sim \mathsf{N}(0, I_d)$. As mentioned earlier, several MGARCH models were proposed in the literature. In our study we focus on two of them: Dynamic conditional correlation (DCC) proposed in Engle (2002) and the more recent Dynamic equicorrelation model (DECO), see Engle and Kelly (2012).

The main feature of the DCC model is the decomposition of H_t into the correlation matrix and two diagonal matrices of variances of the components of ξ_t:

$$\begin{aligned} H_t &= D_t R_t D_t, \text{ where} \\ D_t^2 &= \operatorname{diag}(\sigma_{1,t}^2, \ldots, \sigma_{d,t}^2), \\ R_t &= \operatorname{diag}\left(Q_t^{-1/2}\right) Q_t \operatorname{diag}\left(Q_t^{-1/2}\right), \\ Q_t &= (1-a-b)\overline{Q} + a(\varepsilon_{t-1}\varepsilon_{t-1}^\top) + bQ_{t-1}, \\ \varepsilon_t &\stackrel{\text{def}}{=} D_t^{-1}\xi_t \end{aligned} \tag{4.5}$$

and $\overline{Q}$ is the unconditional covariance matrix of ε_t, $a \geq 0$, $b \geq 0$, $a + b < 1$. The dynamics of the individual series variances $\sigma_{1,t}^2, \ldots, \sigma_{d,t}^2$ has to be specified separately, we assume each of them to follow univariate GARCH(1,1):

$$\sigma_{\ell,t}^2 = \omega_\ell + \alpha_\ell \sigma_{\ell,t-1}^2 + \beta_\ell \zeta_{\ell,t-1}^2, \tag{4.6}$$

where $\zeta_t = \sigma_t z_t$, $z_t \sim \mathsf{N}(0,1)$, $\omega_\ell > 0$, $\alpha_\ell \geq 0$, $\beta_\ell \geq 0$ and $\alpha_\ell + \beta_\ell < 1$.

The DECO model is based on DCC though the two models are not nested. DECO applies the same decomposition of H_t as DCC, but the correlation matrix R_t^{DECO} is a transformation of R_t^{DCC}:

$$R_t^{DECO} = \begin{pmatrix} 1 & \rho_t & \ldots & \rho_t \\ \rho_t & 1 & \ldots & \rho_t \\ \vdots & \vdots & \ddots & \vdots \\ \rho_t & \rho_t & \ldots & 1 \end{pmatrix}, \tag{4.7}$$

where $\rho_t = 1/\{d(d-1)\}\left(\iota^\top R_t^{DCC}\iota - d\right)$ and ι is a unit vector with length equal to d.

As can be easily seen, all non-diagonal elements of R_t^{DECO} are equal. In the following,

this specification will be referred to as "standard DECO".

Another specification of DECO also described in Engle and Kelly (2012) is a so-called Block DECO which allows for a non-homogeneous structure of the vector ξ_t. Similar to the Archimedian copula that can allow for various degrees of strength of the (nonlinear) dependence between different data series, Block DECO can accommodate different intragroup correlations within certain groups of series as well as different intergroup correlations. In general case, if the series can be broken down into V groups of sizes n_1, n_2,...,n_V respectively, the correlation matrix R_t will contain V diagonal blocks with the sides equal to n_1, n_2,...,n_V respectively, and $V(V-1)/2$ off-diagonal blocks:

$$R_t = \begin{pmatrix} (1-\rho_{1,1,t})I_{n_1} & 0 & \dots \\ 0 & \ddots & 0 \\ \vdots & 0 & (1-\rho_{V,V,t})I_{n_V} \end{pmatrix} + \begin{pmatrix} \rho_{1,1,t}J_{n_1} & \rho_{1,2,t}J_{n_1\times n_2} & \dots \\ \rho_{2,1,t}J_{n_2\times n_1} & \ddots & \\ \vdots & & \rho_{V,V,t}J_{n_V} \end{pmatrix}, \quad (4.8)$$

where $\rho_{l,m,t} = \rho_{m,l,t}$ for all l, m.

The correlations in the blocks on the main diagonal are calculated as:

$$\rho_{l,l,t} = \frac{1}{n_l(n_l-1)} \sum_{i\in l, j\in l, i\neq j} \frac{q_{i,j,t}}{\sqrt{q_{i,i,t}q_{j,j,t}}}, \quad (4.9)$$

and the correlations in the off-diagonal blocks are equal to:

$$\rho_{l,m,t} = \frac{1}{n_l n_m} \sum_{i\in l, j\in m} \frac{q_{i,j,t}}{\sqrt{q_{i,i,t}q_{j,j,t}}}, \quad (4.10)$$

$q_{i,j,t}$ is the i,j-th element in the matrix R_t^{DCC}, so Block DECO correlations are in fact average correlations of each block of the matrix R_t in the DCC model.

To denote the structure of Block DECO we will use the same notation that was introduced for the copula structure. Each expression in brackets will refer to one of the V groups of variables.

4.3 Estimation of copula-based multivariate GARCH

As explained above, C-MGARCH models assume that the distribution of the error terms e_t is not limited to standard normal. Instead, e_t is calculated as $e_t = \Sigma^{-1/2}\eta_t$, where η_t is a random vector, $\mathsf{E}(\eta_t|\mathcal{F}_{t-1}) = 0$, $\mathsf{E}(\eta_t\eta_t^\top|\mathcal{F}_{t-1}) = \Sigma$, whose components are in general non-linearly dependent. The exact form of this dependence is described by a copula: $\eta_t|\mathcal{F}_{t-1} \sim F_{1,\dots,d}(\eta_t,\theta) \stackrel{\text{def}}{=} C\{F_1(\eta_{1,t}),\dots,F_d(\eta_{d,t});\theta\}$, where C is a copula, θ is its parameter vector which is assumed constant and $F_1(\eta_{1,t}),\dots,F_d(\eta_{d,t})$ are marginal distributions followed by $\eta_{1,t},\dots,\eta_{d,t}$. In the following we assume $F_1(\eta_{1,t}),\dots,F_d(\eta_{d,t})$

to be standard normal distributions. Since the copula parameter vector θ is assumed constant, the multivariate distribution of η_t is not affected by any past information $\mathcal{F}_{t-1}$.

The assumptions on the distribution of η_t give rise to the following log-likelihood function which should be maximised for model parameter estimation (see Lee and Long (2009)):

$$\mathcal{L}(\varphi, \theta) = \sum_{t=1}^{N} \left\{ \log f_{1,\dots,d}(\eta_t; \theta) + \log \left| \Sigma^{1/2}(\theta) H_t^{-1/2}(\varphi) \right| \right\}, \tag{4.11}$$

where N is the number of observations of vector ξ_t, φ is the vector of the MGARCH parameters defined as $\varphi = (\omega_1, \dots, \omega_d, \alpha_1, \dots, \alpha_d, \beta_1, \dots, \beta_d, a, b)^\top$ for both DCC and DECO models, $f_{1,\dots,d}(\eta_t; \theta)$ is the pdf of η_t. The log-likelihood function should be maximised by choosing the values of φ and θ:

$$(\widehat{\varphi}, \widehat{\theta}) = \arg \max_{\varphi, \theta} \mathcal{L}(\varphi, \theta). \tag{4.12}$$

Taking into account the assumption on the distribution of η_t, we can also write the log-likelihood function for copula-based models as:

$$\mathcal{L}(\varphi, \theta) = \sum_{t=1}^{N} \left[\sum_{i=1}^{d} \{\log f_i(\eta_{i,t})\} + \log c \{F_1(\eta_{1,t}), \dots, F_d(\eta_{d,t}); \theta\} + \log \left| \Sigma^{1/2}(\theta) H_t^{-1/2}(\varphi) \right| \right], \tag{4.13}$$

where $f_i(\eta_{i,t})$ is the marginal pdf of $\eta_{i,t}$ assumed to follow standard normal distribution, $c(\cdot)$ is the copula density function calculated as:

$$c(u_1, \dots, u_d) = \frac{\partial^d C(u_1, \dots, u_d)}{\partial u_1 \dots \partial u_d}. \tag{4.14}$$

Within one MGARCH type (e.g. DCC, or VC), HAC-GARCH represents the most general model with AC-GARCH being its special case which in turn nests standard MGARCH. Since the latter assumes independence of $\eta_{i,t}$, $C(\cdot)$ takes the form of the product copula in this case, hence $\log c(\cdot)$ turns to zero and $\Sigma^{1/2}$ is a unity matrix.

In the more general case of AC-GARCH and HAC-GARCH, matrix Σ has to be normalised for identification ($\sigma_{ii} = 1$). Its off-diagonal elements can be evaluated using Hoeffding's lemma (Hoeffding (1940)):

Hoeffding's Lemma. *Let η_1 and η_2 be random variables with the marginal distributions F_1 and F_2 and the joint distribution F_{12}. If the first and second moments are finite, then*

$$\sigma_{12}(\theta) = \iint_{\mathbb{R}^2} \{F_{12}(\eta_1, \eta_2; \theta) - F_1(\eta_1) F_2(\eta_2)\} \, d\eta_1 d\eta_2. \tag{4.15}$$

Taking Sklar's theorem into account we can rewrite (4.15) in terms of a copula:

$$\sigma_{12}(\theta) = \iint_{\mathbb{R}^2} \left[C \left\{ F_1 \left(\eta_1 \right), F_2 \left(\eta_2 \right); \theta \right\} - F_1 \left(\eta_1 \right) F_2 \left(\eta_2 \right) \right] d\eta_1 d\eta_2. \tag{4.16}$$

5 Simulation study

In order to demonstrate that HAC-MGARCH models can generate time series dynamics that cannot be fully captured by other MGARCH models and to investigate the effects of possible model misspecifications we ran the following Monte-Carlo simulation study using a reduced set of models applied in the empirical study. We pay most attention to the three DCC specifications:

1. Standard four-dimensional DCC.
2. Copula-based four-dimensional DCC with Gumbel AC.
3. Copula-based four-dimensional DCC with Gumbel HAC, $s = ((12)(34))$.

As mentioned in chapter 4, the first two specifications are nested within DCC with Gumbel HAC. For reference purposes we also consider two specifications of DECO:

4. Standard four-dimensional DECO.
5. Block DECO, $s = ((12)(34))$.

A thorough simulation study would require considering multiple parameter sets, applying several copula types and structures and also combining them with all possible Block DECO specifications. However, our aim in this chapter is not to compare the behaviour of all possible specifications, but to investigate on a few representative models if and how simpler specifications can approximate more complex copula-based ones. We assume that more general specifications can generate data dependence structures that cannot be adequately described by models with stricter conditions imposed on errors, e.g. HAC-DCC may not be approximated well enough by AC-DCC, let alone standard DCC. The values of the true parameter vectors are assumed to be reasonably close to the typical empirical parameter vector estimates which is why it is possible to treat the conclusions made in the simulation study as being related to at least a considerable fraction of all practical cases. The structure $s = ((12)(34))$ was considered because it will be used in the empirical study.

The length of each data series is $N = 2000$ observations. The vector of univariate GARCH parameters for the four series is $(\omega_1, \omega_2, \omega_3, \omega_4, \alpha_1, \alpha_2, \alpha_3, \alpha_4, \beta_1, \beta_2, \beta_3, \beta_4)^\top = (0.013, 0.003, 0.014, 0.016, 0.15, 0.2, 0.06, 0.2, 0.6, 0.7, 0.8, 0.6)^\top$. Parameters in the DCC

part are $a = 0.1$, $b = 0.7$. The chosen parameter values correspond to the typical situation encountered in practice: α and a estimates are usually relatively close to their lower boundary 0, whereas β and b estimates are often close to their higher boundary 1. Copula parameters are $\theta = 3$ for AC-DCC and $\theta \stackrel{\text{def}}{=} (\theta_1, \theta_2, \theta_3)^\top = (1.5, 3, 5)^\top$ for HAC-DCC, where θ_1 is the parameter of the outer copula and θ_2 and θ_3 are the parameters of the two subordinated copulas, as in the right panel of Figure 4.1. Nesting conditions $\theta_1 \leq \theta_2$ and $\theta_1 \leq \theta_3$ given in McNeil (2008) are satisfied.

Each of the five models was simulated $J = 1000$ times and each simulated dataset was used to estimate all five specifications. The only fit quality criterion we used is the conditional Kullback-Leibler information criterion (Kullback and Leibler (1951)).

$$KLIC = \mathsf{E}_{\psi_t} \left\{ \log \frac{\psi_t(\xi_t, \varphi, \theta)}{\widehat{\psi}^y(\xi_t, \widehat{\varphi}^y, \widehat{\theta}^y)} \right\}, \tag{5.1}$$

where $\psi_t(\xi_t, \varphi, \theta)$ is the true conditional density of the observation ξ_t calculated using the true parameter vector values $(\varphi, \theta)^\top$ and $\widehat{\psi}_t^y(\xi_t, \widehat{\varphi}^y, \widehat{\theta}^y)$ is the estimate of the conditional density of the observation ξ_t according to the model y calculated using the estimated parameter vector values $(\widehat{\varphi}^y, \widehat{\theta}^y)^\top$.

Expectation in (5.1) is taken with respect to the true conditional density $\psi_t(\cdot|\mathcal{F}_{t-1})$. Since the actual conditional probability corresponds to the true conditional probability, we can estimate the KLIC by taking a simple average of the logarithms of the actual calculated conditional probability ratios. This is why the KLIC estimate takes the form of the difference between the log-likelihood of the true model with true parameters and that of the assumed model with estimated parameters (i.e. the parameters that maximise the respective log-likelihood function), divided by the number of observations:

$$\widehat{KLIC} = \frac{1}{N} \left\{ \mathcal{L}(\varphi, \theta) - \widehat{\mathcal{L}}^y(\widehat{\varphi}^y, \widehat{\theta}^y) \right\}, \tag{5.2}$$

where $N = 2000$ is the length of the simulated series, $\mathcal{L}(\varphi, \theta)$ is the log-likelihood function value according to the true model calculated using true parameter vector values $(\varphi, \theta)^\top$, $\widehat{\mathcal{L}}^y(\widehat{\varphi}^y, \widehat{\theta}^y)$ is the maximum value of the log-likelihood function obtained when estimating model y and $(\widehat{\varphi}^y, \widehat{\theta}^y)^\top$ is the vector that maximises this log-likelihood function. Log-likelihood is calculated as shown in chapter 4.

Empirical kernel density of KLIC of each of the five assumed specification for each of the five true models is shown in Figures 5.1-5.5. In order to estimate kernel density, normal kernel smoother was used with bandwidth h equal to:

$$h = \left(\frac{4}{3J} \right)^{\frac{1}{5}} \text{MAD}_{KLIC}, \tag{5.3}$$

where $\text{MAD}_{KLIC} = 1.4826\,\text{med}\left(|KLIC_j - \text{med}(KLIC)|\right)$ is the median absolute deviation of KLIC, a robust and consistent estimator of its standard deviation (see e.g. Rousseeuw and Croux (1993)), $\text{med}(\cdot)$ is the sample median of $(\cdot)$ and $j = 1, \ldots, J$, where J is the KLIC sample size equal to the number of simulations of each model, i.e. 1000. Bandwidth was estimated for each KLIC distribution separately.

In this formulation of KLIC, its smaller absolute values correspond to a model that is closer to the true one. Estimation of a true specification should result in KLIC estimates distributed closely around zero.

Figure 5.1 shows the distribution of the estimated KLIC based on 5 specifications with standard DCC being the true model. It is clear that the lines corresponding to standard DCC, DCC with Gumbel AC and DCC with Gumbel HAC are indistinguishable from each other since standard DCC is a special case of both AC-DCC and HAC-DCC. The lines corresponding to both DECO specifications are positioned relatively far away from those corresponding to all the DCC specifications which can be interpreted that the ability of DECO to approximate DCC models is rather low.

Figures 5.2 and 5.3 show the results of estimation of the dataset generated by DCC with Gumbel AC and DCC with Gumbel HAC respectively. They represent a similar picture with respect to the position of the kernel density graphs of DECO models. It is also obvious that simpler DCC models fail to capture the dependence structure generated by more general models: standard DCC fails to capture even the dependence structure generated by AC-DCC, let alone HAC-DCC which cannot be satisfactorily approximated by AC-DCC itself.

Standard and Block DECO are the models that generated the datasets analysed in Figures 5.4 and 5.5 respectively. These figures are given for reference purposes since one could not expect DCC models to perfectly approximate DECO models since both model classes are non-nested. However, one can see from Figure 5.4 that all three lines demonstrating the distribution of the KLIC estimate obtained from the DCC models estimation is positioned relatively closely to that of the standard DECO compared with the situation in Figures 5.1, 5.2 and 5.3 where all KLIC estimate distributions resulting from the DECO models estimation are quite remote from those corresponding to the true models. In Figure 5.5, the KLIC estimate distributions resulting from the true Block DECO and all DCC models estimation are again reasonably close to each other with the distribution of the KLIC estimate related to the standard DECO model being distant from that of the true model showing that standard DECO fails to approximate Block DECO satisfactorily. The ideas conveyed by the KLIC estimate kernel density figures may be summarised in the following way. First, the KLIC estimate corresponding to the true specification is indeed always distributed around zero and has a relatively small variance. Second and most important, the KLIC estimate distribution resulting

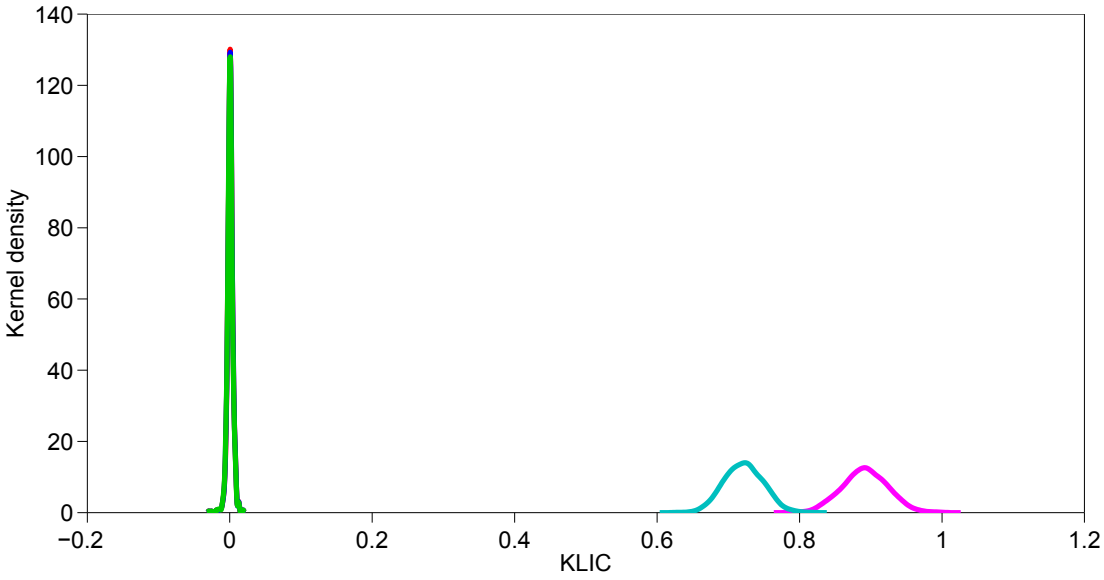

Figure 5.1: KLIC kernel density estimation under various assumed MGARCH specifications, true model: standard DCC. Assumed specifications: standard DCC, DCC with Gumbel AC, DCC with Gumbel HAC, $s = ((12)(34))$(last three lines coincide), standard DECO, Block DECO, $s = ((12)(34))$.

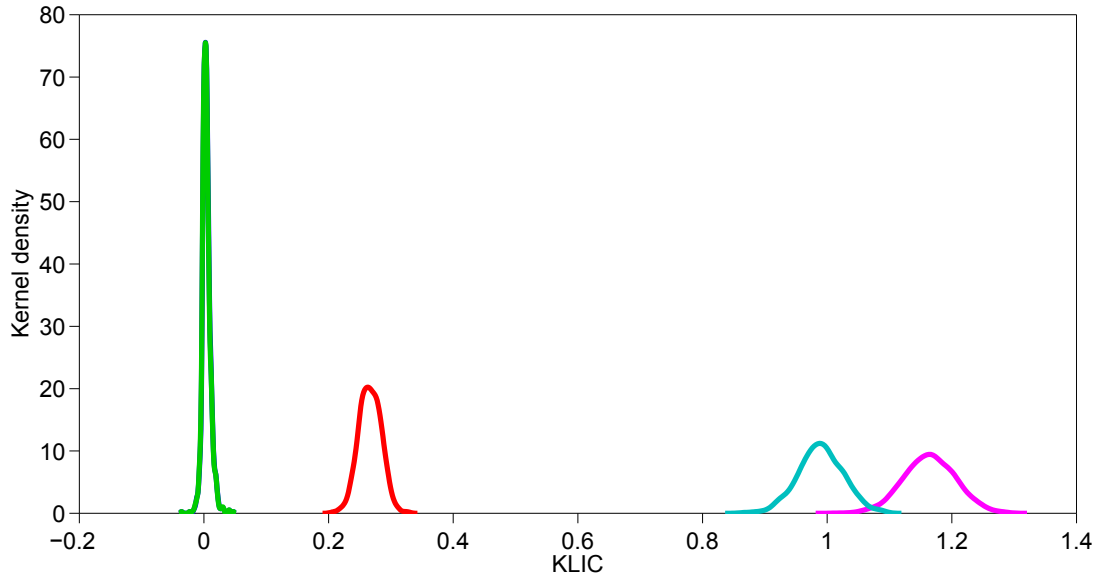

Figure 5.2: KLIC kernel density estimation under various assumed MGARCH specifications, true model: DCC with Gumbel AC. Assumed specifications: standard DCC, DCC with Gumbel AC, DCC with Gumbel HAC, $s = ((12)(34))$(last two lines coincide), standard DECO, Block DECO, $s = ((12)(34))$.

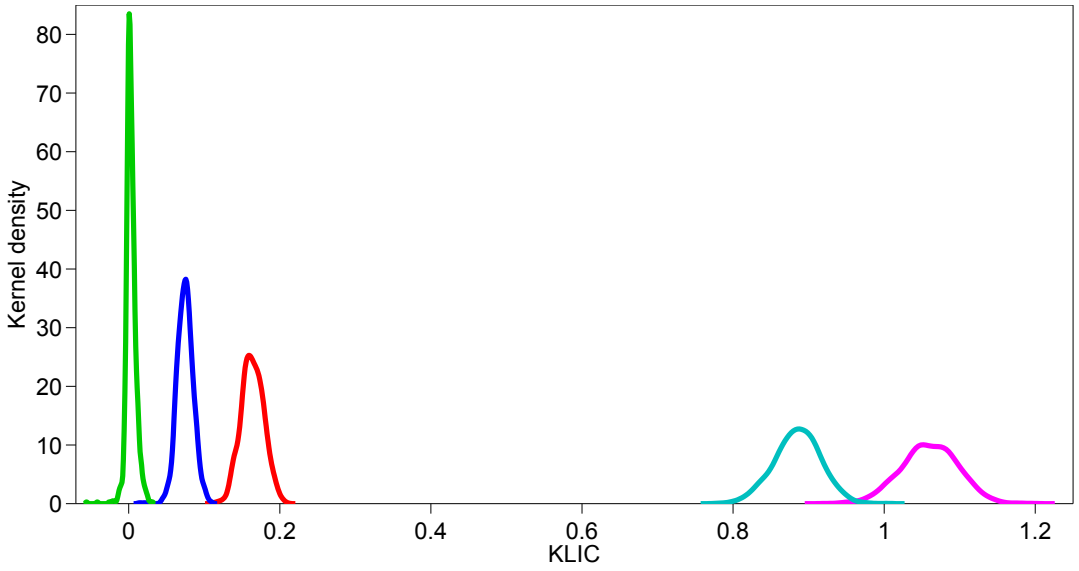

Figure 5.3: KLIC kernel density estimation under various assumed MGARCH specifications, true model: DCC with Gumbel HAC, s=((12)(34)). Assumed specifications: standard DCC, DCC with Gumbel AC, DCC with Gumbel HAC, $s = ((12)(34))$, standard DECO, Block DECO, $s = ((12)(34))$.

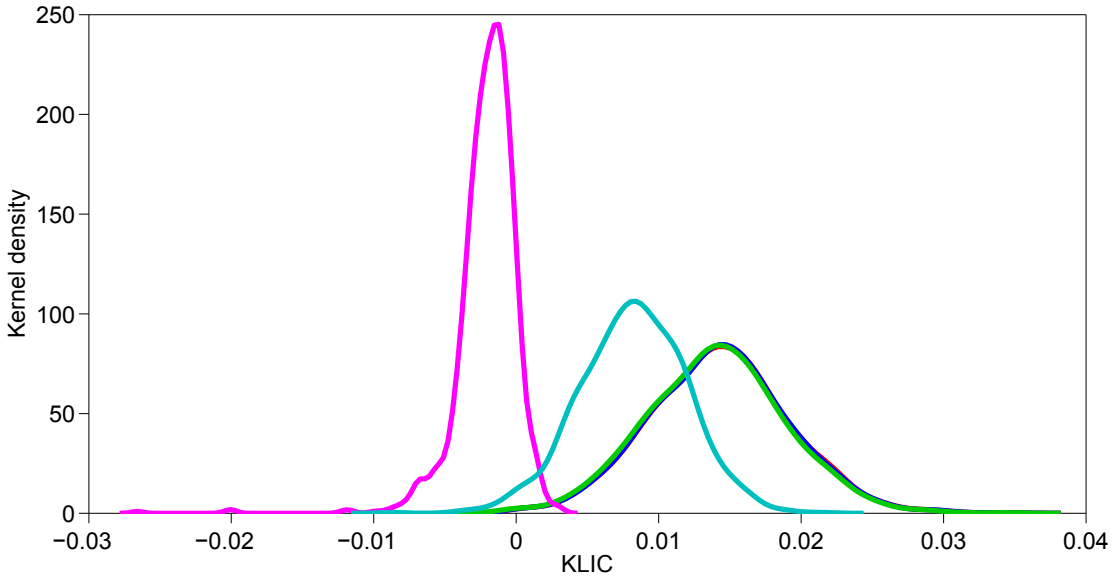

Figure 5.4: KLIC kernel density estimation under various assumed MGARCH specifications, true model: standard DECO. Assumed specifications: standard DCC, DCC with Gumbel AC, DCC with Gumbel HAC, $s = ((12)(34))$(last three lines almost coincide), standard DECO, Block DECO, $s = ((12)(34))$.

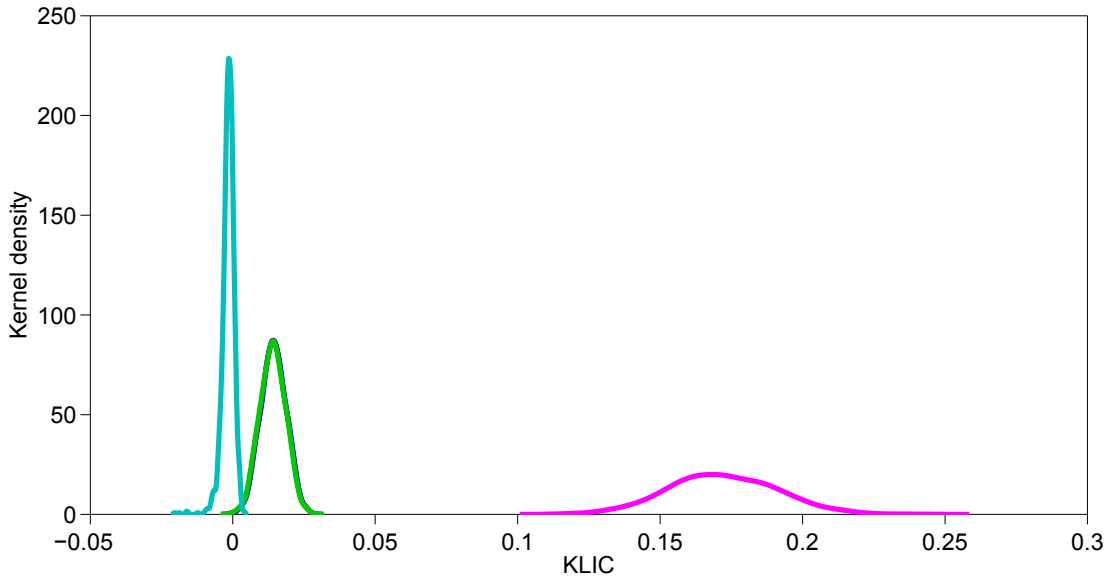

Figure 5.5: KLIC kernel density estimation under various assumed MGARCH specifications, true model: Block DECO, $s = ((12)(34))$. Assumed specifications: standard DCC, DCC with Gumbel AC, DCC with Gumbel HAC, $s = ((12)(34))$(last three lines almost coincide), standard DECO, Block DECO, $s = ((12)(34))$.

from HAC-DCC estimation obviously matches the KLIC estimate distribution resulting from the estimation of less sophisticated nested models (standard DCC, AC-DCC), whereas the opposite is not true: i.e. KLIC estimate distributions under various assumed specifications hardly intersect when the true model is HAC-DCC which means that a HAC-DCC dataset will be adequately described only by estimating a HAC-DCC specification. Third, DCC models seem to be able to approximate dependence structures generated by DECO models better than the latter can approximate the dependence structures generated by the former. This implies that our suggestion in chapter 4 is supported: HAC-DCC specification in fact looks promising since it can capture dependence structures that other models cannot, at the same time it can approximate dependence structures defined by other models, albeit with different quality.

6 Empirical study

In the empirical section we conduct a multi-stage analysis of the futures prices data with the goal to compare the ability of the mathematical tools described in chapters 3 and 4, namely of the two-factor model of the futures curve and recent developments in the MGARCH modelling to capture and forecast the joint dynamics of the two commodity futures curves. We simulated 1000 random portfolios and computed value-at-risk forecasts based on the assumptions and estimated parameters of 11 models that were estimated on 495 different time windows and conducted the value-at-risk backtesting procedure. Our main criterion in defining best-performing models was the average precision of value-at-risk estimation.

In particular, first the two-factor model of the futures curve (Ohana (2010)) is calibrated which allows to estimate the long- and short-term shocks series for both commodities. In line with the two-factor model, the deterministic and the random parts of the shocks series are analysed separately. The subsequent analysis is performed on a rolling window of 500 trading days with the step equal to 5 trading days. On each window the vector autoregression (VAR) model is estimated in order to extract the deterministic part from the shocks series, then 11 C-MGARCH specifications are estimated on the residual series of the VAR model. Combined with the VAR forecast, each of the 11 C-MGARCH models produces a forecast of the joint distribution of the four daily shocks for each of the 5 days following the last day of each estimation period. This forecast is transformed to the forecast of the joint distribution of the daily return of a particular portfolio of futures contracts for these 5 days. The forecast of the joint distribution of the portfolio return is used to predict daily value-at-risk for a particular day and particular predefined exceedance probability for this portfolio.

In this chapter, first the data is presented, then the two-factor model of the futures curve is calibrated, after that basic properties of the shocks series are investigated. It is followed by the presentation of the VAR model and description of the C-MGARCH specifications used in the study. Finally, the value-at-risk backtesting procedure is explained in detail, and its results are discussed.

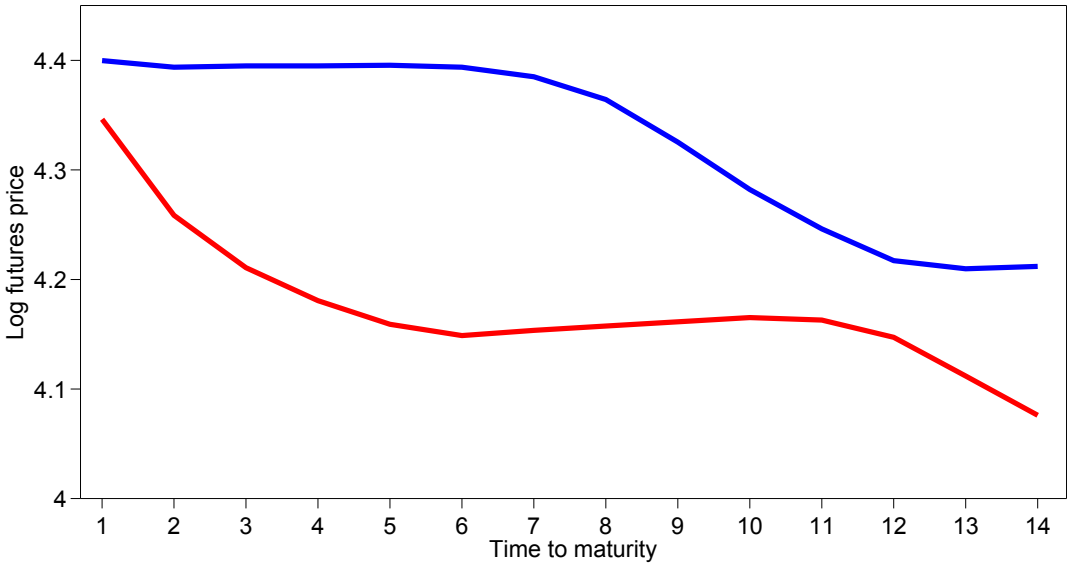

Figure 6.1: Logarithmised heating oil futures curves on 01.02.2000 and 22.06.2000 in US cents per gallon.

6.1 Data description

We use the same, but a longer, dataset as in Ohana (2010), i.e. NYMEX daily heating oil and natural gas futures prices from 1.02.2000 to 19.12.2011 (exact names of the contracts are New York Harbour No. 2 Heating Oil Futures and Henry Hub Natural Gas Futures). Similar to Ohana (2010), we also consider first 14 maturities for both commodities. The data was obtained from Bloomberg. Figure 6.1 demonstrates a heating oil futures curve on 1.02.2000 and 22.06.2000. On both dates the curve exhibits two obvious features. The first one is backwardation, i.e. a situation where more distant maturities trade at lower prices than less distant ones which is a normal situation on the commodity markets (see e.g. Gabillon, 1991) and reflects (buyers') "preference for the present time whatever the reasons are". The other one is seasonality: contracts expiring in winter are generally more expensive than those expiring in summer. It is also obvious that because of the backwardation, rising slopes are much less pronounced than falling ones. As was shown in Chapter 3, the seasonality feature can be disregarded for the purposes of our analysis.

	Heating oil	Natural gas
k	3.10	2.86
Percentage of explained variance	98.23	93.80

Table 6.1: Characteristic numbers of the commodities and explanatory power of the two-factor model of the futures curve.

6.2 Two-factor model calibration

The first step of the analysis is the calibration of the two-factor model described in chapter 3. $N = 2975$ daily observations of futures price returns for $M = 14$ maturities were used. The calibration results, i.e. the values of the parameter k for both commodities as well as the quality of the model fit are given in Table 6.1. The model explains over 90% of the return variance for both commodities which is a good result. For heating oil futures returns, the model even explains over 98% of the variance. Due to a longer time span used, our results are slightly different from those obtained by Ohana (2010), i.e. we find the evidence for even higher explanatory power of the model for the heating oil futures data and a little lower explanatory power for the natural gas data. Both calibrated parameters are somewhat lower than in Ohana's study. Calibration of the two-factor model allows to estimate the four shocks series that are shown in Figure 6.4. Their properties are discussed in the next section.

6.3 Basic properties of the shocks, levels and slopes time series

As mentioned above, the shocks series will remain in the focus of the subsequent analysis. One of the main properties of any time series that should be checked before any further econometric analysis is stationarity. Figure 6.4 clearly suggests that both long- and short-term shocks are concentrated around zero and do not seem to have any trend. Since all four shocks are "components" of a return of a futures contract, they are analogous to stock market returns that are common to assume to be stationary (Tsay (2002)). This is why it is reasonable to assume shocks to be stationary, too.

Figure 6.2: Heating oil and natural gas levels (1.02.2000 – 19.12.2011).

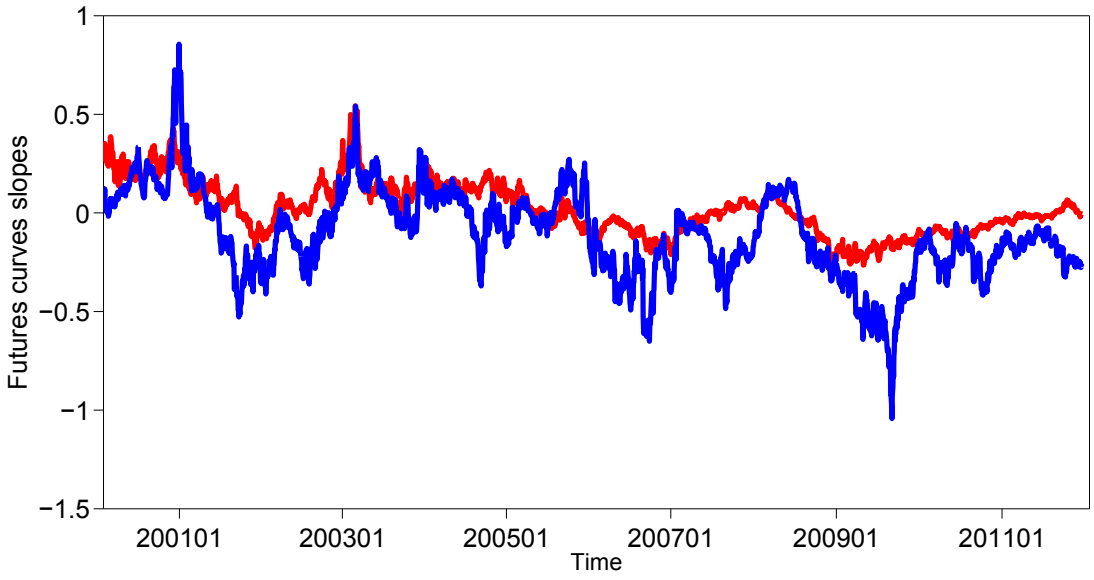

Figure 6.3: Heating oil and natural gas slopes (1.02.2000 – 19.12.2011).

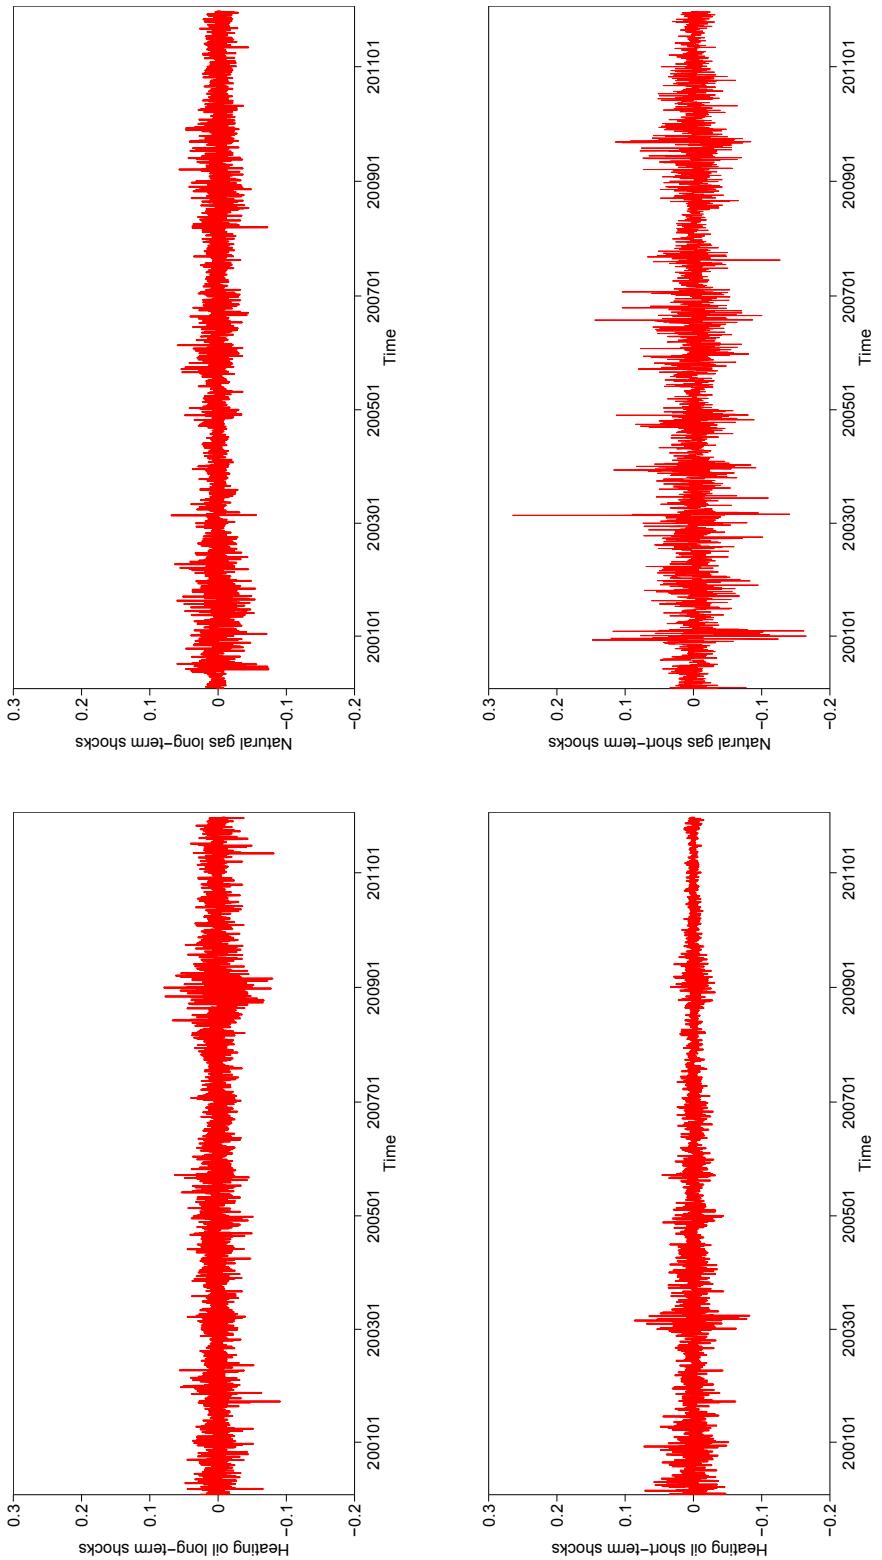

Figure 6.4: Estimated shocks in the two-factor model of the futures curve: heating oil long-term shocks (upper left), natural gas long-term shocks (upper right), heating oil short-term shocks (lower left), natural gas short-term shocks (lower right) (1.02.2000 – 19.12.2011).

Series	ADF (p-value)	KPSS (p-value)
Heating oil levels	-1.43 (0.55)	4.82 (0.01)
Natural gas levels	-1.66 (0.45)	1.96 (0.01)
Heating oil slopes	-3.07 (0.03)	2.84 (0.01)
Natural gas slopes	-3.71 (0.00)	1.85 (0.01)
Heating oil long term shocks	-19.80 (0.00)	0.17 (0.10)
Natural gas long term shocks	-19.03 (0.00)	1.05 (0.01)
Heating oil short term shocks	-19.91 (0.00)	0.15 (0.10)
Natural gas short term shocks	-20.34 (0.00)	0.05 (0.10)

Table 6.2: Results of the stationarity tests of the shocks series, p-values in brackets.

Slopes and levels represent two different kinds of transformation of the shocks series. Given the nature of these series, intuitively we cannot expect levels to be stationary because global factors affecting both supply and demand for commodities can change in any unpredictable way causing deterministic or stochastic trends in levels development. This intuition is partially supported by Figure 6.2 which demonstrates that the levels of both commodities during the period in consideration were experiencing an overall rising at least up to 2008. Slopes represent a weighted sum of presumably stationary time series. But since the weights of the summands change constantly, it is hard to predict whether the slopes series would be stationary or not. The intuition tells us that the slopes may have natural limits and therefore are expected to have a bouncing nature rather than any kind of trend. However, Figure 6.3 provides some evidence in favour of a possible negative trend in both slopes series during the period in consideration.

To further investigate unit-root properties of all 8 aforementioned series, Augmented Dickey-Fuller (ADF) test in its "only constant" specification and the "no trend" specification of the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test were carried out. The results are reported in Table 6.2. Qualitatively similar results were obtained under other test specifications. As can be concluded from Table 6.2, ADF test does not reject the null hypothesis of unit root for the levels series, but does reject it for all the other series. The KPSS test rejects the stationarity hypothesis for both levels and slopes series, but does not reject it for the shocks series. These results are in line with our expectations. We will also use slopes as regressors in the model. The values of the Partial Autocorrelation Function (PACF) of the shocks series (Figure 6.5) suggest that estimation of autoregressive models can yield useful information about the shocks dynamics, but the order of autoregressive models is not expected to be high which will be confirmed during the VAR model estimation.

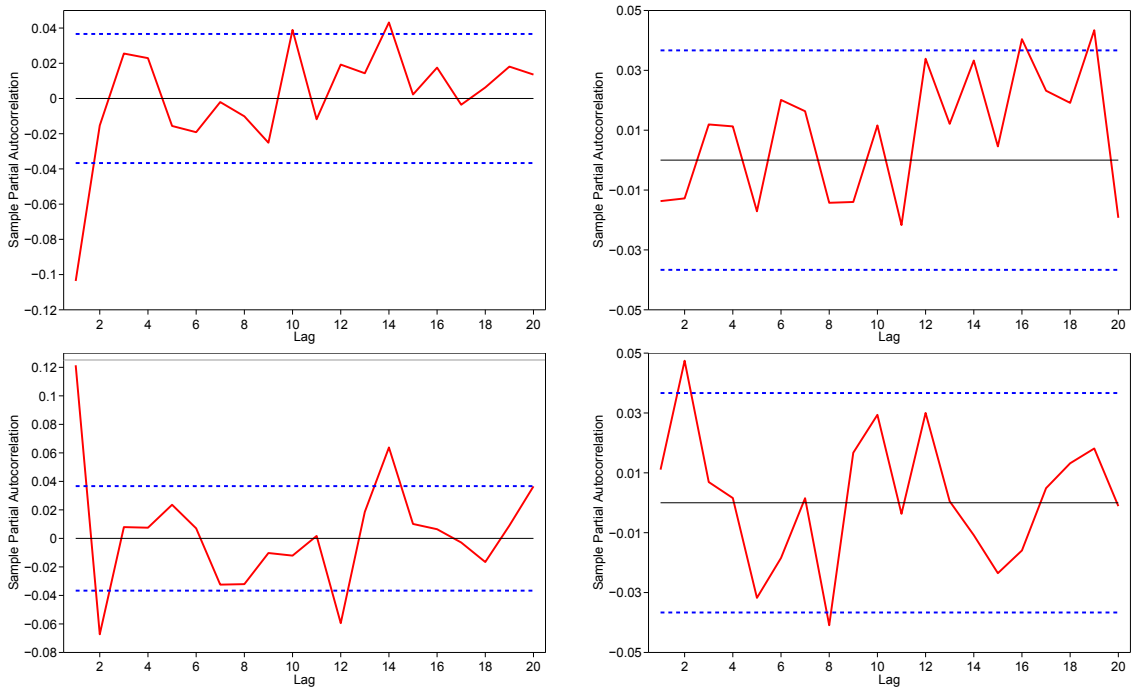

Figure 6.5: Partial autocorrelation function (PACF) for heating oil long-term shocks (upper left), natural gas long-term shocks (upper right), heating oil short-term shocks (lower left), natural gas short-term shocks (lower right).

6.4 Vector autoregression

After both long- and short-term shocks are calculated, the next step is to find out if any of the past available information can be used to predict their dynamics. This may mean e.g. estimation of autoregressive models as suggested in the previous section. The approach applied in Ohana (2010) is to estimate a VECM (vector error correction model). According to its formulation used in Ohana (2010), the levels of the two commodities are cointegrated. Following the approach of Engle and Granger (1987), the long-term relationship between the commodity levels is estimated in a separate model and turns out to be well approximated by a piecewise linear function. In the next step, the residuals in this model, i.e. deviations from the long-term relationship are used as an exogenous variable in the VAR model. In this paper we do not apply the VECM framework for two reasons. First, even with cointegration accounted for, the model in Ohana (2010) explains a too small proportion of the shocks' variance: the R-squared of the four equations within the VECM model explaining the dynamics of the heating oil long-term shocks, natural gas long-term shocks, heating oil short-term shocks and natural gas short-term shocks was only 3.19%, 1.89%, 2.21% and 1.72% respectively. Second, there is enough evidence of the fact that the cointegration link between oil and gas was broken over the course of 2009 due to marked changes in the US gas market.

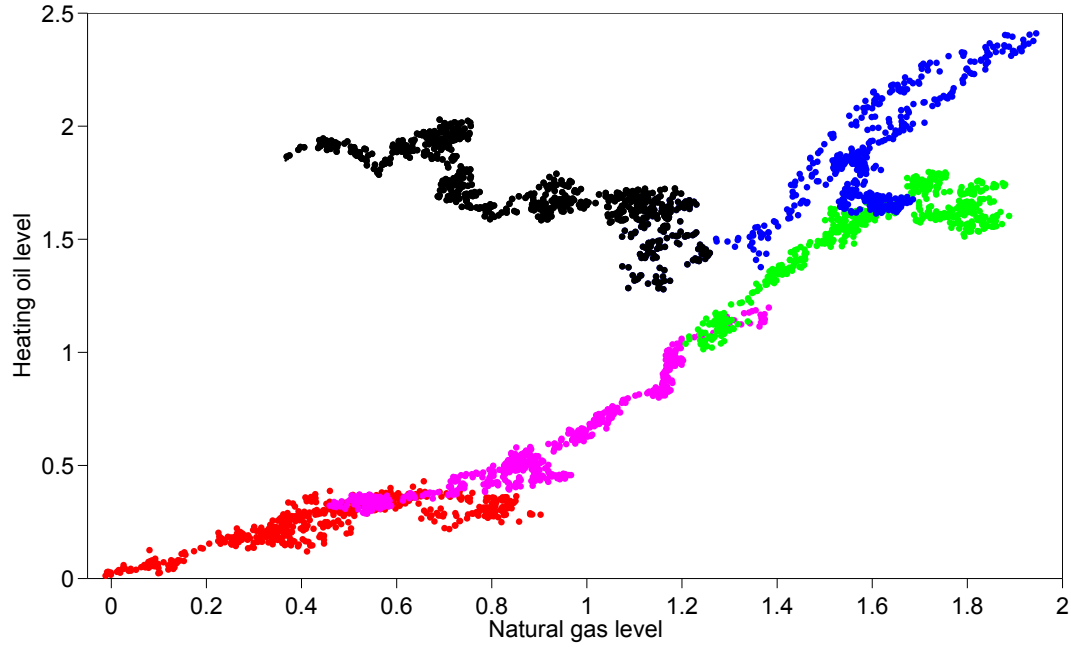

Figure 6.6: Long-term relationship between heating oil and natural gas levels (20000201 - 20020621 , 20020622 - 20041108, 20041109 - 20070329, 20070330 - 20090120, 20090121 - 20111219)

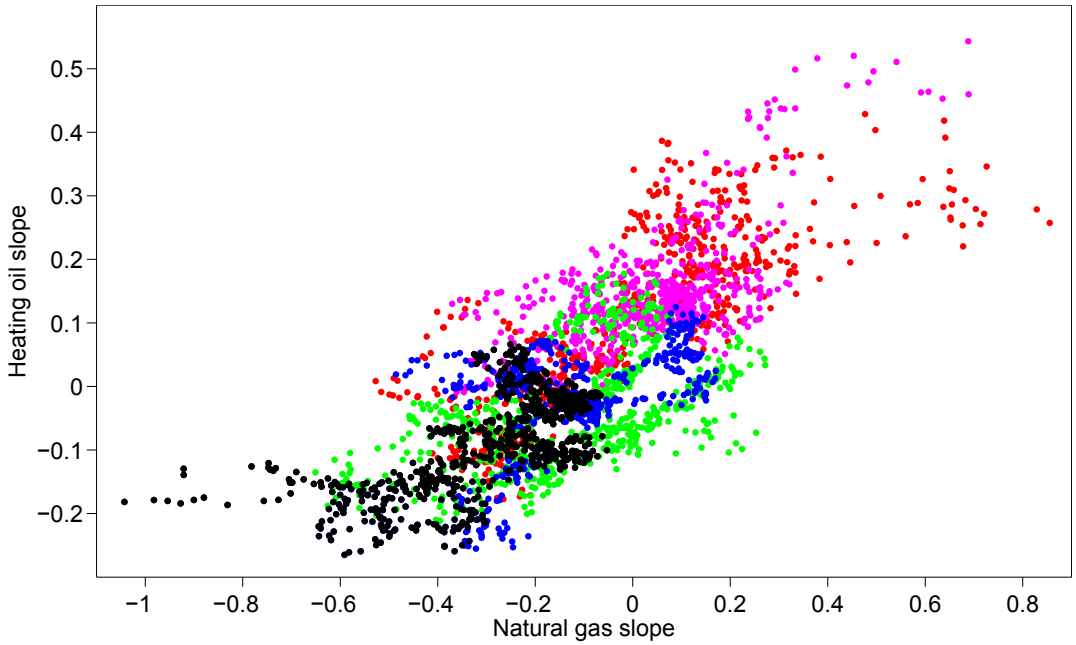

Figure 6.7: Long-term relationship between heating oil and natural gas slopes (20000201 - 20020621 , 20020622 - 20041108, 20041109 - 20070329, 20070330 - 20090120, 20090121 - 20111219)

A quick look at Figure 6.6 confirms this: whereas oil prices have recovered after their slump during the 2008-2009 crisis, gas prices have not. De Bock and Gijón (2011) cite additional supply of non-conventional gas (above all shale gas) and high storage levels in the US market as main reasons behind the relative weakness of US gas prices and the loosening of the link between WTI and natural gas prices. This period can also be denoted as a period of increased uncertainty in the gas market because it appears to be difficult to reliably estimate the amount of the gas that can be extracted, and serious corrections of the available reserves in the future cannot be excluded. It is also reasonable to assume that the most the part of 2009 and 2010 constitute a transition to a new long-term relationship between the two commodities, because there seems to be little reason for the common factors on the demand side and for the link between the two commodities to disappear completely. As a supporting argument, one could refer to the Annual Energy Outlook prepared by the US Energy Information Administration (EIA). In this report, among other things they provide their view on the future ratio of low-sulfur light crude oil price to Henry Hub natural gas price which in their reference scenario appears to have reached its new long-term level around 2010-2011 after a significant rise during 2007-2010. This projection can also refer to the ratio of heating oil price to natural gas price since heating oil was shown to be cointegrated with WTI, a low-sulfur light grade of oil and one of the world benchmarks in oil pricing (see e.g. Hartley et al. (2008)). Figure 6.6 shows that during the period covered by Ohana's (2010) study, the cointegration approach seemed reasonable, but the newer data, shown in black, give some impression of negative correlation between heating oil and natural gas price levels. Considering this negative relationship as a long-term dependence cannot have any theoretical justification. This is why the VECM structure seemed to make sense when only the data roughly up to February 2009 (the dataset used in Ohana (2010)) were available, but not afterwards.

As a result, the model specification used in this paper is a vector autoregression (VAR) with the maximum lag of 1 and two extra regressors - lagged slopes of both futures curves:

$$Z_t = \mu + \Gamma Z_{t-1} + \Pi A_{t-1} + \xi_t, \tag{6.1}$$

where $Z_t = \left(\delta_t^{oil,L}, \delta_t^{gas,L}, \delta_t^{oil,S}, \delta_t^{gas,S}\right)^\top$ is the shocks vector, μ is a vector of constants, $A_t = \left(X_t^{oil}, X_t^{gas}\right)^\top$ is the slopes vector, $\mathsf{E}(\xi_t|\mathcal{F}_{t-1}) = 0$, ξ_t is explained by C-MGARCH, Γ and Π are parameter matrices.

The maximum lag of 1 was chosen based on Hannan-Quinn and Bayesian Schwartz information criteria that both were minimal for 1 lag for all 495 time windows with only one exception. For both criteria their generalisation for multivariate processes was calculated, see Lütkepohl and Krätzig (2004) for details.

The model is estimated on 495 time windows of size equal to 500 trading days. The

first time window begins on the first observation day, 1.02.2000, and each subsequent window begins five days later than the previous one. We treat VAR estimation as a procedure with the only aim to extract the deterministic component from the vector Z_t. The exact form of the autoregressive relation is not as important to us as the obtained residual series which is why no VAR results are reported here.

6.5 C-MGARCH specifications

C-MGARCH models are used in our study to describe the variance-covariance dynamics of the random components ξ_t of the shocks series. Residuals in the VAR model described in the previous section are the estimates of the random shocks components and will be used as inputs in the C-MGARCH models.

In chapter 4, we mentioned that HAC represents a flexible way to describe various dependence structures which in theory should lead to better performance of HAC-MGARCH models as compared to that of AC-MGARCH or standard MGARCH models. We check this hypothesis by comparing forecasting properties of standard DCC, AC-DCC and HAC-DCC. We also compare their performance with that of standard and Block DECO. By considering DECO specifications, we did not pursue the goal to compare DCC and DECO per se, but we note that the situation where non-copula based DECO models perform better than copula-based DCC would be undesirable because depending on the comparative performance of non-copula based DCC and DECO it would suggest that either the choice of DCC as the basic specification was wrong or that the copula assumption does not help to better capture the data dynamics structure. As shown later in the backtesting section, on average such a situation is not the case.

Normally, estimation of any HAC-MGARCH model based on a copula with a predetermined generator implies not only estimation of parameters, but also the choice of an optimal structure. Here we do not apply any methods to estimate the copula structure, but we assume the following two structures to have enough economic motivation: $s_1 = ((ol\ gl)(os\ gs))$ and $s_2 = ((ol\ os)(gl\ gs))$, where ol, gl, os, gs are the components of the vector η_t in the C-MGARCH model (see section 4.3) corresponding to heating oil long-term shocks, natural gas long-term shocks, heating oil short-term shocks and natural gas short-term shocks respectively. Indeed, the shocks can be naturally classified into two groups in two different ways, and for each case one can assume that the shocks in both pairs are more closely related to each other than to the shocks of the other pair. We do not have any reason to treat either structure as more probable, which is why each structure gives rise to a separate HAC-DCC model. Moreover, we impose the same two structures on the vector ξ_t in the Block DECO model.

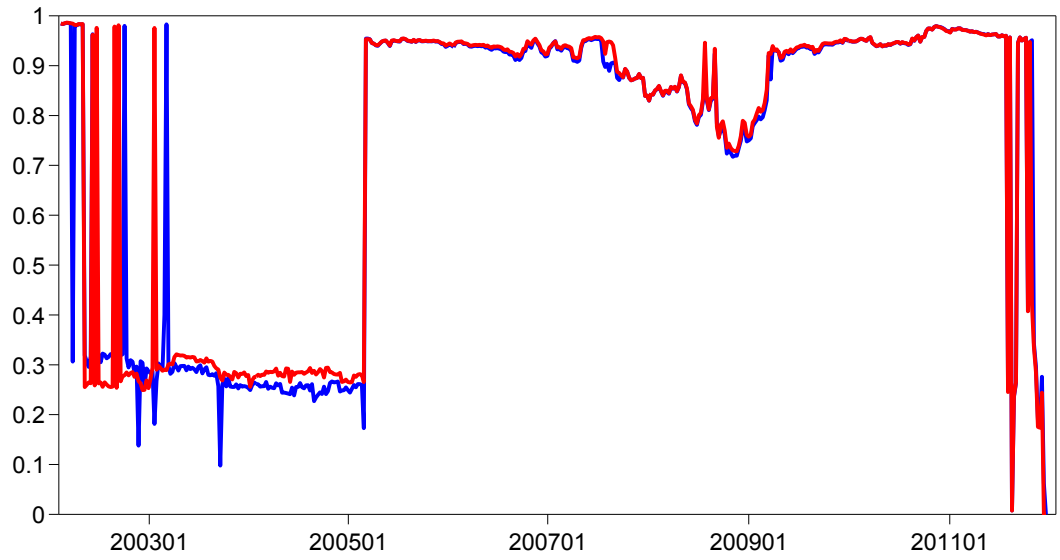

Figure 6.8: Parameter b in the DCC part: standard DCC, DCC with Gumbel HAC, (s_2) (1.02.2000 – 19.12.2011). Values were estimated during the 500-day period ending on the respective day.

All in all, we have 7 DCC modifications: standard DCC and copula based DCC with Gumbel AC, Clayton AC, Gumbel HAC (s_1), Clayton HAC (s_1), Gumbel HAC (s_2) and Clayton HAC (s_2). Three DECO specifications estimated for reference purposes are: standard DECO, Block DECO with s_1 and s_2. We also estimate univariate GARCH which is effectively a special case of C-MGARCH and can be characterised as DCC or DECO with constant zero correlation between the components of ξ_t.

Figures 6.8 and 6.9 demonstrate how parameter b in the main DCC equation and Kendall's τ (calculated as $\tau = (\theta - 1)/\theta$ for the Gumbel copula) corresponding to the three copula parameters in DCC with Gumbel HAC, s_2 vary from period to period which can be seen as a justification to perform model estimation on several time windows.

6.6 Portfolio value-at-risk backtesting

Portfolio value-at-risk backtesting is a simple and useful tool for quality assessment of a model that belongs to the GARCH family since value-at-risk is a risk measure widely used in practice. The aim of this procedure is to compare the precision of value-at-risk forecasts produced by different models for different futures portfolios. It is easy to show that according to the two-factor model of the futures curve (Ohana (2010)) the return of any portfolio of the futures contracts can be expressed as a "portfolio" of the four

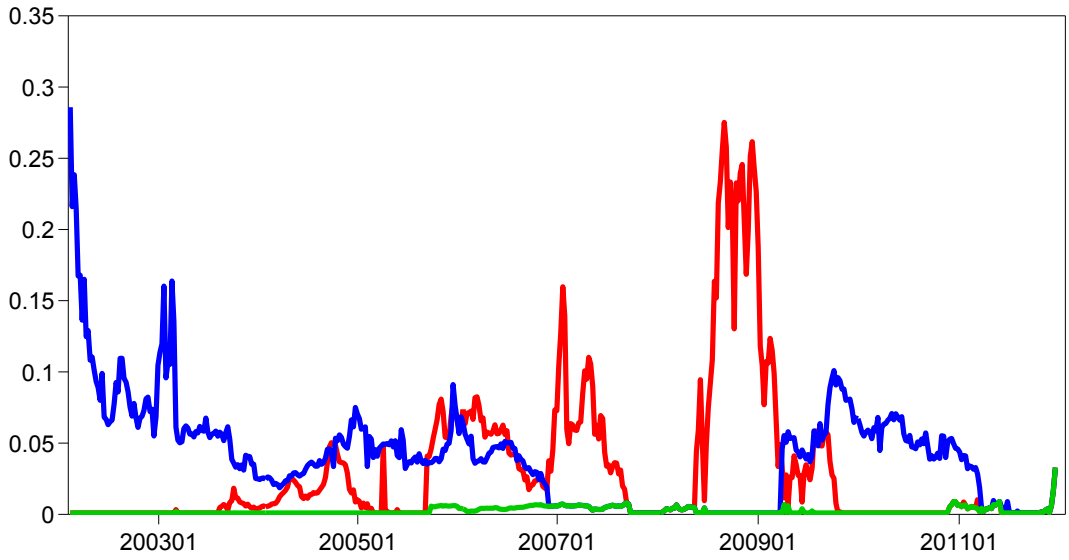

Figure 6.9: Kendall's τ corresponding to the copula parameters in DCC with Gumbel HAC, s_2: parameter of the subcopula linking heating oil long- and short-term shocks residuals (red), parameter of the subcopula linking natural gas long- and short-term shocks residuals (blue) and parameter of the outer copula (green) (1.02.2000 – 19.12.2011). Values were estimated during the 500-day period ending on the respective day

shocks. This means that if the joint distribution of the shocks can be estimated for a particular day, the distribution for any futures portfolio return, i.e. for any weighted sum of the shocks can be easily obtained.

If the portfolio return forecast for an arbitrary day $t+1$, Ret_{t+1} is calculated as $Ret_{t+1} = w_s^\top Z_{t+1}$, where $w_s = (w_{s1}, w_{s2}, w_{s3}, w_{s4})^\top$ is the vector of the shocks weights corresponding to the vector of the futures weights and Z_{t+1} is the future vector of shocks, then value-at-risk (V@R) at level $0 < \alpha < 1$ for day $t+1$ is defined as $V@R_{t+1}(\alpha) \stackrel{\text{def}}{=} F^{-1}_{Ret_{t+1}}(\alpha)$. Keeping in mind that every shock δ_t is composed of a deterministic component λ_t explained by the VAR model and a zero-expectation random component ξ_t whose variance-covariance matrix is explained and can be predicted by a C-MGARCH model, the expression for value-at-risk takes the form:

$$V@R_{t+1} = w_s^\top \lambda_{t+1} + F^{-1}_{w_s^\top \xi_{t+1}}(\alpha). \tag{6.2}$$

With estimated VAR parameters, it is straightforward to calculate the forecast for λ_{t+1} because all necessary information up to the time t is available. For standard MGARCH models, the expression $F^{-1}_{w_s^\top \xi_{t+1}}(\alpha)$, where $\xi_{t+1} = H^{1/2}_{t+1} e_{t+1}$ is evaluated based on the as-

sumption on standard normal distribution of e_{t+1} which implies that $w_s^\top \xi_{t+1}$ is also distributed normally: $w_s^\top \xi_{t+1} \sim \mathsf{N}(0, w_s^\top H_{t+1} w_s)$. For AC-MGARCH and HAC-MGARCH such estimation is not possible since $e_{t+1} = \Sigma^{-1/2}\eta_{t+1}$ is not normal, hence the whole distribution of $H_{t+1}^{1/2} e_{t+1}$ has to be simulated based on the estimated copula parameters governing the distribution of η_{t+1}. For each distribution 3000 points are simulated. After that it is easy to obtain the distribution of any weighted sum of the four components of $H_{t+1}^{1/2} e_{t+1}$ and to estimate the empirical quantile $F^{-1}_{w_s^\top \xi_{t+1}}(\alpha)$.

We estimated both VAR and each of the 11 C-MGARCH specifications on 495 time windows with a size equal to 500 observations. Taking into account that each window begins five days later than the previous one, having estimated the parameters on a particular window, we used them to calculated a forecast for the 5 days following the last day of the window. For each 5-day forecast only the parameters were taken constant, the information on futures returns was updated every day. The whole time period covered in the study contains 2975 daily observations. Since the first 500 were used for the model estimation only, forecasts are available for $2975 - 500 = 2475$ days.

For a particular futures portfolio it is useful to estimate the exceedance rate $\widehat{\alpha}$ of every model which is the share of the observations for which the actual portfolio return is lower that the corresponding value-at-risk forecast:

$$\widehat{\alpha}_w \stackrel{\text{def}}{=} n^{-1} \sum_{t=1}^{n} \mathbf{I}\left\{ Ret_t < \widehat{V@R}_t(\alpha) \right\}, \tag{6.3}$$

where $w = (w_1, \ldots, w_{28})^\top$ are the futures portfolio weights, $n = 2475$ is the number of forecast values for each portfolio, and the relative deviation of the exceedance rate from the true α is:

$$d_w \stackrel{\text{def}}{=} \frac{\widehat{\alpha}_w - \alpha}{\alpha}.$$

A model predicts value-at-risk perfectly if $\widehat{\alpha}$ and α coincide, and $d_w = 0$. Another procedure to measure the accuracy of the value-at-risk forecast is the unconditional coverage test of Kupiec (1995) with the null hypothesis H_0: $\widehat{\alpha} = \alpha$. The likelihood ratio test statistic takes the form:

$$LR_{Kupiec} = 2 \log \left\{ \left(\frac{1 - \widehat{\alpha}}{1 - \alpha} \right)^{n - \mathbf{I}(\alpha)} \left(\frac{\widehat{\alpha}}{\alpha} \right)^{\mathbf{I}(\alpha)} \right\}, \tag{6.4}$$

where $\mathbf{I}(\alpha) = \sum_{t=1}^{n} \mathbf{I}\left\{ Ret_t < \widehat{V@R}_t(\alpha) \right\}$. The test statistic follows $\chi^2(1)$ under H_0.

We generated a set W of $|W| = 1000$ portfolios $w^1, \ldots, w^{1000}$ of $p = 28$ components with weights of each component in each portfolio $w_1^q, \ldots, w_p^q$, $q = 1, \ldots, |W|$ uniformly distributed over the simplex $S^p = \{(y_1, \ldots, y_p) | \sum_{i=1}^{p} y_i = 1\}$. For weights generation we used a variation of procedure 4 described in the review of weights generation algorithms

by Wang and Zionts (2006). Procedure 4 was shown to efficiently generate uniformly distributed portfolio weights. However, by design it can produce only positive weights distributed over the simplex $S^p = \{(y_1, \ldots, y_p) | y_i \geq 0, \sum_{i=1}^{p} y_i = 1\}$, whereas in our case it is also desirable to allow for negative weights which would correspond to taking short positions in contracts. Unlike short positions in stocks, short positions in futures do not cause any additional transaction costs as compared to long positions. This is an argument in support of equal treatment of long and short positions both of which are likely to be taken e.g. by a trader seeking arbitrage opportunities arising from futures pricing deviating too far from some theoretical relationship. Moreover, higher variance of the weights will allow us to check if a particular model can predict not only the left tail of the shocks distribution correctly, but also its other intervals. Indeed, testing a model on a portfolio will all negative weights (an extreme situation that cannot happen in our case) can tell us if the model can capture the right tail of the shocks distributions correctly. This is why the weights were allowed to take negative values, too. The exact procedure we followed to generate both positive and negative weights is as follows:

1. Generate a set Λ_1 of $|\Lambda_1| = 1000$ portfolios of $p = 28$ components with weights of each component in each portfolio $\lambda_1^{1,q}, \ldots, \lambda_p^{1,q}$, $q = 1, \ldots, |\Lambda_1|$ uniformly distributed over the simplex $S^p = \{(y_1, \ldots, y_p) | y_i \geq 0, \sum_{i=1}^{p} y_i = 1\}$. Multiply all weights in all portfolios by 2.

2. Generate a set Λ_2 of $|\Lambda_2| = 1000$ portfolios of $p = 28$ components with weights of each component in each portfolio $\lambda_1^{2,q}, \ldots, \lambda_p^{2,q}$, $q = 1, \ldots, |\Lambda_1|$ uniformly distributed over the simplex $S^p = \{(y_1, \ldots, y_p) | y_i \geq 0, \sum_{i=1}^{p} y_i = 1\}$.

3. Let W be a set of $|W| = 1000$ portfolios with weights calculated as $w_i^q = \lambda_i^{1,q} - \lambda_i^{2,q}$ for all $i = 1, \ldots, p$ and $q = 1, \ldots, |W|$.

4. The mean weight $\overline{w}$ in the resulting portfolios set W is equal to $\overline{w} = 1/28 = 0.0357$ and the weight's standard deviation $\sigma_w = 0.077$ (measured across all contracts and portfolios).

We replace one of the simulated portfolio by an equally-weighted portfolio, $w_i^1 = 1/28 = 0.0357$. For the evaluation of the model performance on the whole portfolio set W we will use the average exceedance rates and average p-values of the Kupiec test measured across all $|W|$ portfolios. Additionally we calculate the average relative deviation of d_w and its standard error:

$$A_W = \frac{1}{|W|} \sum_{w \in W} d_w, \qquad D_W = \left\{ \frac{1}{|W|} \sum_{w \in W} (d_w - A_W)^2 \right\}^{1/2}. \tag{6.5}$$

Table 6.3 shows the results of the value-at-risk backtesting for the equally-weighted

portfolio of 28 contracts and Table 6.4 summarises the results of the value-at-risk backtesting for all 1000 portfolios.

As can be seen from Table 6.3, for $\alpha = 10\%$ and $\alpha = 5\%$, H_0 of the Kupiec test is not rejected at the 5% significance level for all models except univariate GARCH which implies that treating shocks series as independent leads to severe underestimation of possible losses. However, best performers for $\alpha = 10\%$ and $\alpha = 5\%$ are Block DECO with s_2 and standard DECO respectively. We also see that copula-based DCC models do not perform better than standard DCC for $\alpha = 5\%$ (though they do for $\alpha = 10\%$). Only $\alpha = 1\%$ is the case where a HAC-based model is a clear leader: DCC with Clayton HAC, s_1 shows the best result. Moreover, this model is one of the only two for which the hypothesis $\widehat{\alpha} = 1\%$ is not rejected at the 5% significance level. It is necessary to emphasise that while best performing models were able to forecast value-at-risk fairly well for $\alpha = 10\%$ and $\alpha = 5\%$ (though on average all models have a relatively low exceedance rate, i.e. they overestimate losses), the value-at-risk forecasts for $\alpha = 1\%$ exhibit a relatively high exceedance rate which means that the fat-tail distributions of the portfolio returns are not fully captured by the considered models. Good performance of DCC with Clayton HAC, s_1 for $\alpha = 1\%$ can be explained by the fact that the Clayton copula describes lower-tail dependence. This is also why exceedance rates of the models based on the Clayton copula are mostly lower than those of the corresponding models based on the Gumbel copula. Overall results of the value-at-risk backtesting for the equally-weighted portfolio cannot be seen as satisfactory with regard to the comparably better performance that we could expect from HAC-based models.

However, the situation is absolutely different on the aggregate level, i.e. for 1000 simulated portfolios on average. For all three α values, HAC-based DCC models are best in class. They do not only produce the most accurate value-at-risk forecasts on average as demonstrated by their exceedance rates that are closest to the respective required α values, but also generate these accurate forecasts on a regular basis as witnessed by their low A_W and D_W values. It is also a very encouraging result that for two of the three α values, standard DCC performs worse than DECO in most cases. This means that a copula assumption can increase forecast qualities of a poorly-performing model significantly, so that it even overtakes other benchmarks. Some general conclusions, such as lower exceedance rates of DCC models with Clayton copula as compared to the corresponding DCC models with Gumbel copula, generally significantly worse performance of all models for $\alpha = 1\%$ (Kupiec test's null hypothesis is not rejected only two specifications of DCC based on the Clayton copula at the 10% signifiance level) and a better ability of DCC with Clayton HAC, s_1 to capture lower fat tails of the portfolio return distributions, that were drawn from the value-at-risk backtesting for the equally-weighted portfolio, remain valid for the analysis of the 1000 portfolios.

Model	$\alpha = 10\%$	$\alpha = 5\%$	$\alpha = 1\%$
DCC	9.616(0.522)	4.889(0.799)	1.495(0.021)
DCC with Gumbel AC	9.778(0.712)	5.253(0.567)	1.455(0.033)
DCC with Clayton AC	9.657(0.567)	4.848(0.728)	1.455(0.033)
DCC with Gumbel HAC, s_1	9.697(0.614)	5.131(0.765)	1.495(0.021)
DCC with Clayton HAC, s_1	9.657(0.567)	4.687(0.470)	**1.374(0.077)**
DCC with Gumbel HAC, s_2	9.778(0.712)	5.131(0.765)	1.414(0.051)
DCC with Clayton HAC, s_2	9.737(0.662)	4.889(0.799)	1.495(0.021)
Standard DECO	9.778(0.712)	**5.051(0.908)**	1.495(0.021)
Block DECO, s_1	9.697(0.614)	4.848(0.728)	1.455(0.033)
Block DECO, s_2	**10.141(0.815)**	5.333(0.451)	1.616(0.005)
Univariate GARCH	16.283(0.000)	11.071(0.000)	4.162(0.000)

Table 6.3: Value-at-risk backtesting results for the equally-weighted portfolio: $\hat{\alpha}$ (in %) and Kupiec test p-values (in brackets). Results of the models yielding highest p-values are shown in bold.

Value-at-risk backtesting results can be illustrated by depicting value-at-risk for each forecast day along with actual returns. Figure 6.10 provides more insights into the third column ($\alpha = 5\%$) of Table 6.3 and shows such graphs for standard DCC, DCC with Gumbel AC, DCC with Gumbel HAC, s_2 and Block DECO, s_2. In Figure 6.10, actual returns are shown in blue if they were higher than the predicted value-at-risk and in red in the opposite case (exceedance event). The value-at-risk profiles look very similar in these plots, the only difference that is easy to notice is that certain data points are coloured differently in different plots meaning that different models produce somewhat different portfolio return distribution forecasts.

Another perspective of the value-at-risk forecasts can be obtained by plotting their kernel densities over the whole forecasting period for each model and each α. Figures 6.11 and 6.12 show the value-at-risk kernel densities for $\alpha = 5\%$ and $\alpha = 1\%$ respectively for each of the 11 considered models. Kernel densities were evaluated using normal kernel smoother and the optimal bandwidth was estimated as in chapter 5. Only left tails of the resulting kernel density graphs are shown in the figures, but now it is much easier to notice that the models behave differently on average. Models based on the Clayton copula are more inclined to provide low value-at-risk estimates than other models which is demonstrated by the fatter left tails of the respective kernel densities and which is especially pronounced for $\alpha = 1\%$. This is in line with our conlusion from the analysis of Table 6.3. Block DECO, s_2 produces relatively high value-at-risk estimates for $\alpha = 5\%$ and $\alpha = 1\%$ that have a higher chance to be exceeded. In Table 6.3 we see that disregarding univariate GARCH, this specification performs worst and has the highest exceedance rates which is in line with relatively thin left tails of its predicted value-at-risk kernel density in Figures 6.11 and 6.12.

Finally, we plot entire estimated daily portfolio return distributions for three particular days produced by the following models: standard DCC (used as a benchmark) and three best performers in value-at-risk backtesting for the equally-weighted portfolio: Block DECO, s_2, standard DECO and DCC with Clayton HAC, s_1. By doing this, we show that the strength of the nonlinear dependence as expressed by the copula parameters has an impact on the skewness of the resulting distribution and hence on its position relative to those generated by the simpler standard DCC and by DECO models. Figure 6.13 shows the kernel density of the portfolio return distribution for 06.07.2011, the day that was preceded by a period characterised by a relatively high nonlinear dependence, hence relatively high Clayton HAC coefficients. Kernel density was again calculated using normal kernel smoother and the optimal bandwidth was estimated as in chapter 5. It is easy to see that the distribution predicted by DCC with Clayton HAC return distribution has a fatter left tail than the ones produced by the other three models which in turn results in lower value-at-risk predicted by the HAC-DCC model. In Figure 6.14, the distributions predicted by the same models for 18.09.2003 are depicted. The copula parameters estimated during the 500-day period before this day are near average, and the distribution generated by DCC with Clayton HAC is closer to the other three than in Figure 6.13, and its left tail is only marginally fatter than that of generated by the other three models. Finally, the distribution forecast for 16.03.2005, a day for which a low nonlinear dependence is predicted, is shown in Figure 6.15. As a result of the low dependence, DCC with Clayton HAC almost degenerates to standard DCC which is why the kernel densities for these two models almost coincide, and the copula-based model does not exhibit a fatter left tail.

Summing up, one can conclude that as expected, HAC-DCC models can be useful for risk-management purposes. In particular, on average they generate more accurate value-at-risk forecasts for various futures portfolios than nested models with stricter assumptions and even some other benchmark models (DECO). In particular, one model based on the Clayton HAC due to the ability of this copula to capture lower tail dependence, is able to approximate the futures portfolios distribution in the region corresponding to particularly low quantiles (as $\alpha = 1\%$) better than other models. Moreover, unlike for most other models, this approximation is statistically significant as shown by the Kupiec test results.

	$\alpha = 10\%$		$\alpha = 5\%$		$\alpha = 1\%$	
Model	$\overline{\hat{\alpha}}$ (p-value)	$A_W(D_W)$	$\overline{\hat{\alpha}}$ (p-value)	$A_W(D_W)$	$\overline{\hat{\alpha}}$ (p-value)	$A_W(D_W)$
DCC	9.760(0.528)	-0.024(0.045)	5.051(0.476)	0.010(0.087)	1.450(0.068)	0.450(0.127)
DCC with Gumbel AC	9.806(0.557)	-0.019(0.044)	5.146(0.449)	0.029(0.089)	1.471(0.064)	0.471(0.139)
DCC with Clayton AC	9.678(0.495)	-0.032(0.045)	5.029(0.488)	0.006(0.086)	1.400(0.107)	0.400(0.132)
DCC with Gumbel HAC, s_1	9.853(0.571)	-0.015(0.044)	5.147(0.449)	0.029(0.088)	1.440(0.072)	0.440(0.122)
DCC with Clayton HAC, s_1	9.677(0.488)	-0.032(0.046)	4.981(0.475)	-0.004(0.089)	**1.362(0.141)**	**0.362(0.131)**
DCC with Gumbel HAC, s_2	**9.965(0.615)**	**-0.004(0.041)**	5.255(0.470)	0.051(0.073)	1.501(0.039)	0.501(0.116)
DCC with Clayton HAC, s_2	9.800(0.586)	-0.020(0.040)	**5.021(0.519)**	**0.004(0.077)**	1.449(0.062)	0.449(0.113)
Standard DECO	9.961(0.598)	-0.004(0.045)	5.251(0.493)	0.050(0.076)	1.560(0.027)	0.560(0.132)
Block DECO, s_1	9.869(0.572)	-0.013(0.046)	5.149(0.515)	0.030(0.079)	1.481(0.055)	0.481(0.136)
Block DECO, s_2	10.111(0.563)	0.011(0.047)	5.367(0.396)	0.073(0.079)	1.575(0.023)	0.575(0.131)
Univariate GARCH	15.318(0.001)	0.532(0.123)	9.886(0.001)	0.977(0.219)	3.807(0.000)	2.807(0.711)

Table 6.4: Summary of the value-at-risk backtesting results for 1000 portfolios: $\overline{\hat{\alpha}}$ is the average exceedance rates across all portfolios (in %), p-values are average Kupiec test p-value across all portfolios. A_W and D_W are defined as in (6.5). Results of the models yielding highest average p-values are shown in bold.

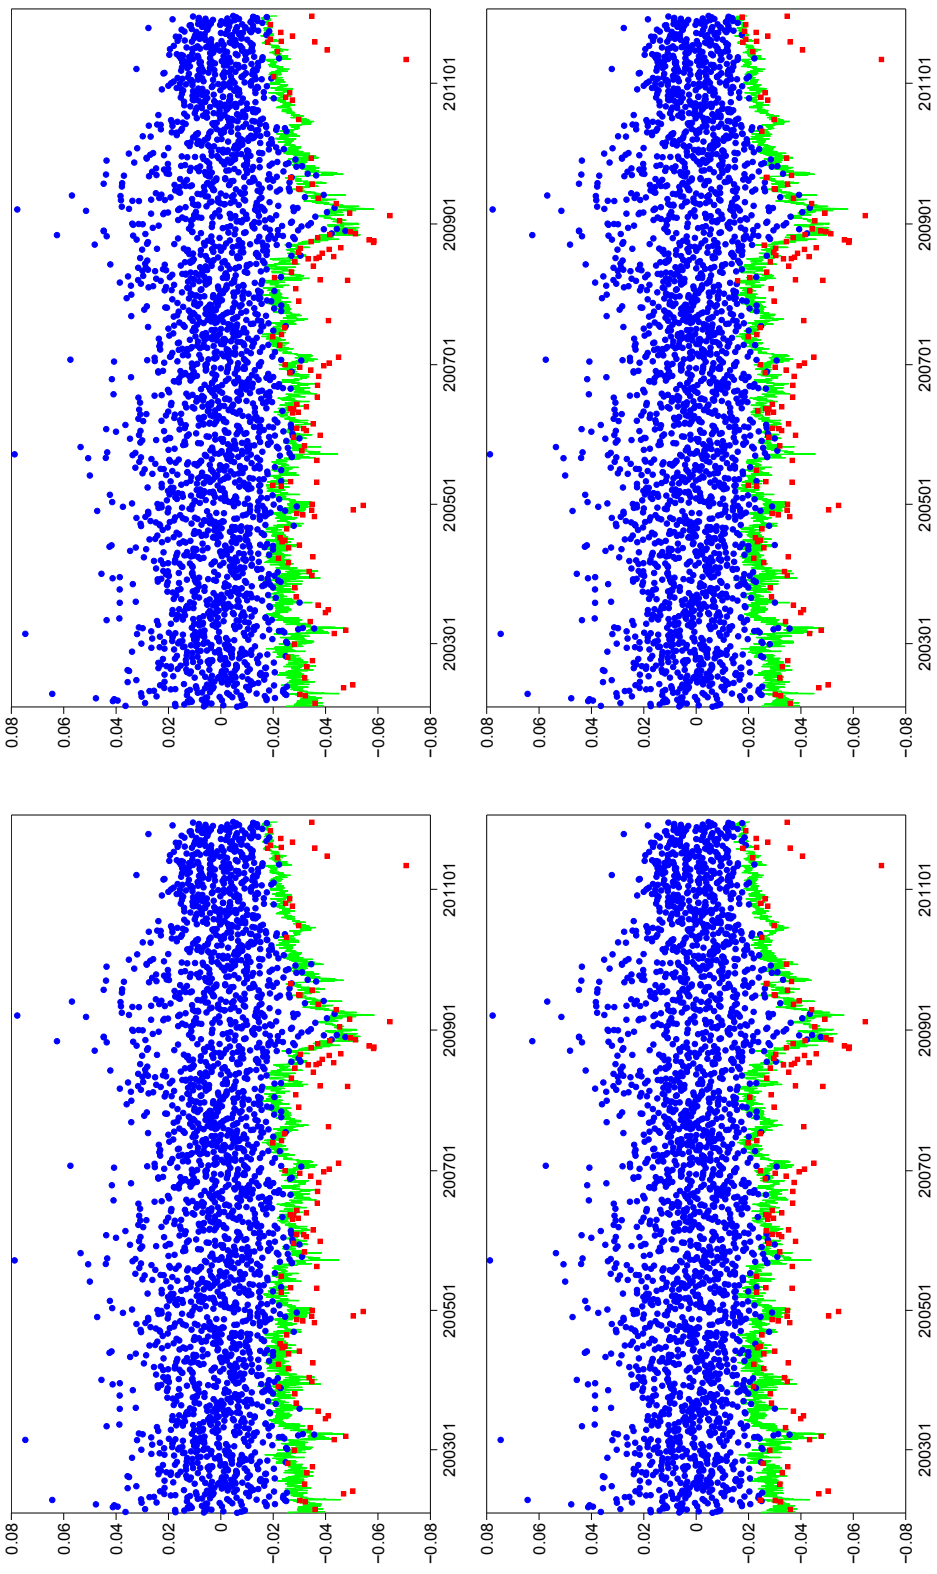

Figure 6.10: Value-at-risk ($\alpha = 5\%$) exceedance plots for standard DCC (upper left), DCC with Gumbel AC (upper right), DCC with Gumbel HAC, s_2 (lower left) and Block DECO, s_2 (lower right) (6.02.2002 – 19.12.2011). Data series: predicted value-at-risk (green line), portfolio returns higher than predicted value-at-risk (blue dots) and portfolio returns lower than predicted value-at-risk (red squares).

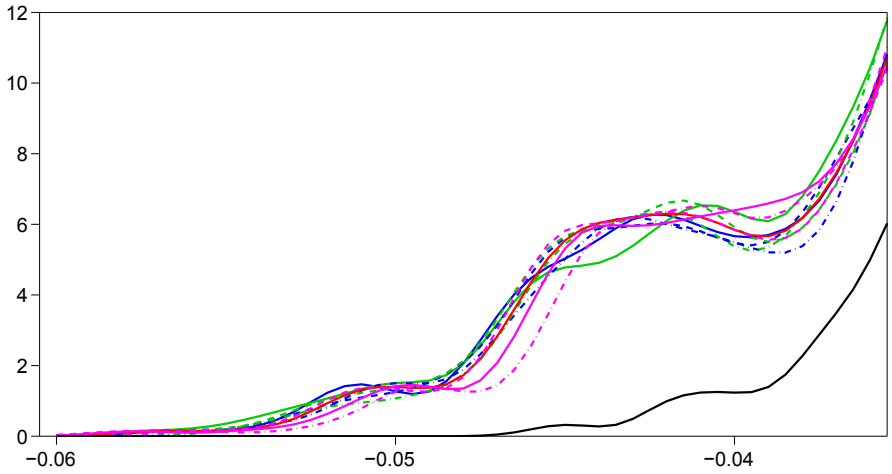

Figure 6.11: Kernel density of the value-at-risk ($\alpha = 0.05$) forecast by standard DCC (solid line), copula-based DCC with Gumbel AC (solid line), Gumbel HAC, s_1 (dashed line), Gumbel HAC, s_2 (dashed-dot line), Clayton AC (solid line), Clayton HAC, s_1 (dashed line), Clayton HAC, s_2 (dashed-dot line), standard DECO (solid line), Block DECO, s_1 (dashed line), Block DECO, s_2 (dashed-dot line), univariate GARCH (solid line).

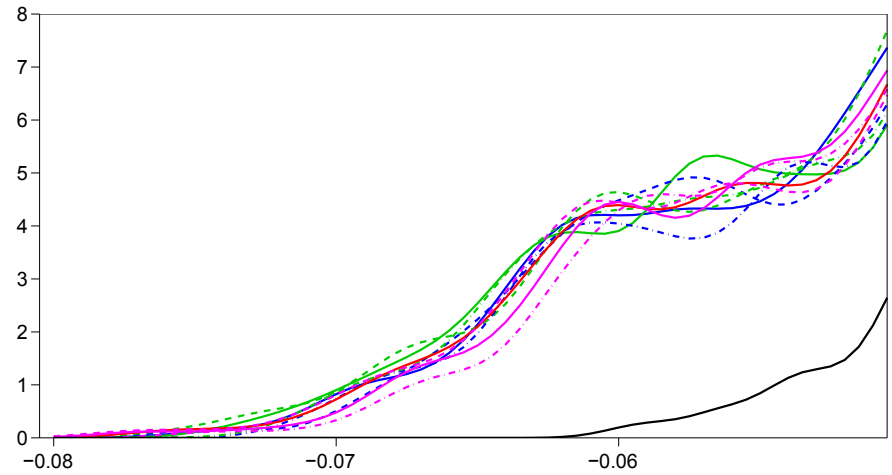

Figure 6.12: Kernel density of the value-at-risk ($\alpha = 0.01$) forecast by standard DCC (solid line), copula-based DCC with Gumbel AC (solid line), Gumbel HAC, s_1 (dashed line), Gumbel HAC, s_2 (dashed-dot line), Clayton AC (solid line), Clayton HAC, s_1 (dashed line), Clayton HAC, s_2 (dashed-dot line), standard DECO (solid line), Block DECO, s_1 (dashed line), Block DECO, s_2 (dashed-dot line), univariate GARCH (solid line).

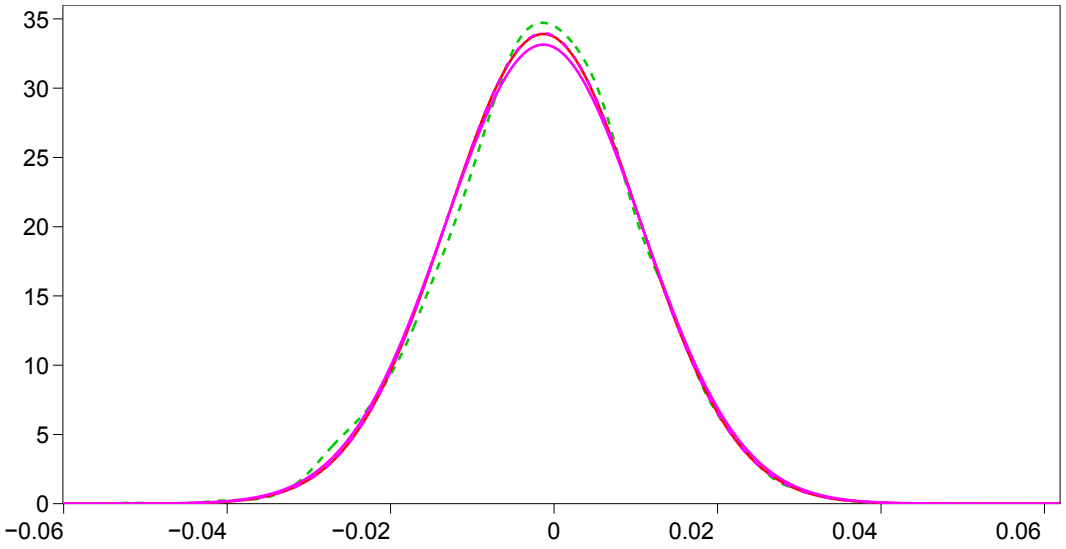

Figure 6.13: Equally-weighted portfolio return distribution forecast for 06.07.2011, estimated copula parameters in DCC with Clayton HAC, s_1: $\theta_1 = 0.002$, $\theta_2 = 0.085$, $\theta_3 = 0.182$ (strong dependence). Standard DCC (solid line), DCC with Clayton HAC, s_1 (dashed line), standard DECO (solid line), Block DECO, s_2 (dashed-dot line).

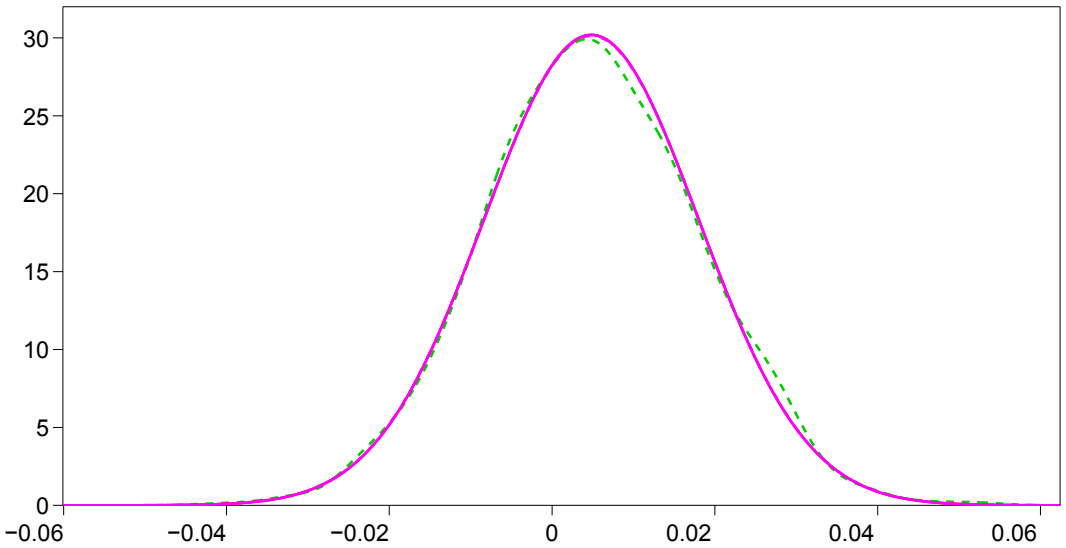

Figure 6.14: Equally-weighted portfolio return distribution forecast for 18.09.2003, estimated copula parameters in DCC with Clayton HAC, s_1: $\theta_1 = 0.028$, $\theta_2 = 0.028$, $\theta_3 = 0.043$ (medium dependence). Standard DCC (solid line), DCC with Clayton HAC, s_1 (dashed line), standard DECO (solid line), Block DECO, s_2 (dashed-dot line).

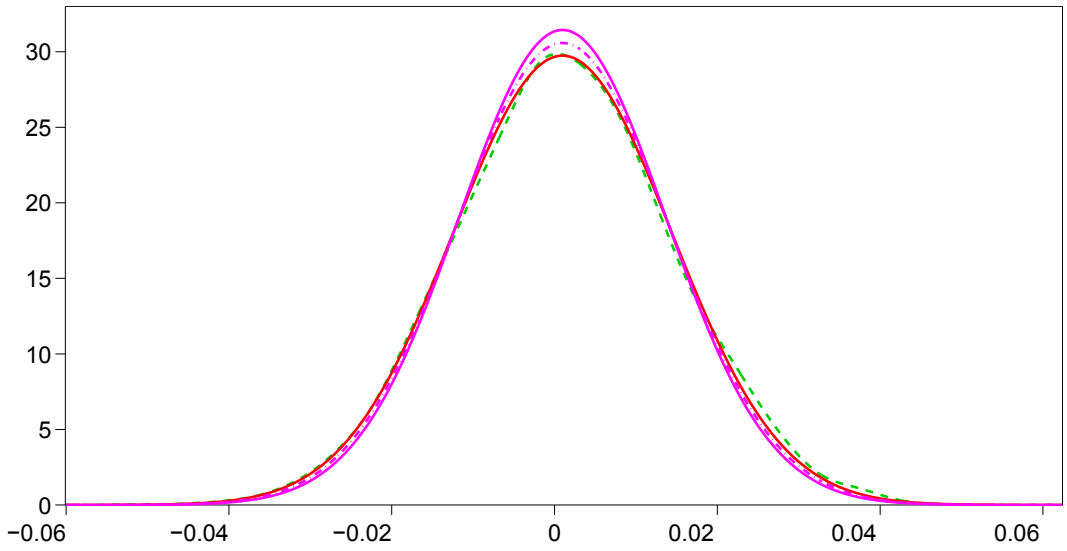

Figure 6.15: Equally-weighted portfolio return distribution for 16.03.2005, estimated copula parameters in DCC with Clayton HAC, s_1: $\theta_1 = 0.002$, $\theta_2 = 0.002$, $\theta_3 = 0.002$ (weak dependence). Standard DCC (solid line), DCC with Clayton HAC, s_1 (dashed line), standard DECO (solid line), Block DECO, s_2 (dashed-dot line).

7 Conclusions

This thesis addresses the issue that can be of significant importance to many agents involved in commodity trading. The study further develops the work of Ohana (2010) and models the dynamics of the heating oil and natural gas futures curves within one model. Multi-stage analysis of a large set of futures prices is carried out: first, futures return series are transformed to the shocks series with the help of the two-factor model of the futures price, then two components of the shocks were analysed separately. For the analysis of the variance-covariance structure of the vector of the random shocks component, HAC-MGARCH models, a recent development in the MGARCH class were used. This analysis allowed to forecast the distribution of the returns of any portfolios composed of the available futures contracts for short time periods. As shown in the study, value-at-risk estimates derived from the forecasts produced by HAC-DCC models are accurate, and these models outperform other benchmark models on a consistent basis as shown by the value-at-risk backtesting procedure carried out on a set of 1000 simulated futures portfolios.

The research can be extended along the following directions. First, it is possible to let MGARCH models account for the possible seasonality in the dynamics of the variance-covariance matrix of the vector of the random shocks components as suggested in Ohana (2010). Second, it is worthwhile to further exploit the advantages of the Archimedian copula class and use other copula types and different structures, in particular a procedure to estimate an optimal structure can be implemented which would be especially helpful if the dimensionality of a problem were to increase with the inclusion of other commodities in the analysis. Moreover, one can further combine copulas with different MGARCH models, e.g. copula-based DECO models could be a promising forecasting tool. Finally, more sophisticated methods to the determination of the estimation time window, such as local adaptive methods, can be applied which can lead to the better understanding of the time evolution of the processes and more precise forecasts.

Bibliography

Annual Energy Outlook. U.S. Energy Information Administration. Report DOE/EIA-0383(2011), April 2011. URL `http://www.eia.gov/forecasts/archive/aeo11/`.

F. Asche, P. Osmundsen, and M. Sandsmark. The UK market for natural gas, oil and electricity: are prices decoupled? *The Energy Journal*, 27(2):27–40, 2006.

L. Bachmeier and J. Griffin. Testing for market integration: crude oil, coal, and natural gas. *The Energy Journal*, 27(2):55–72, 2006.

S. Borak and R. Weron. A semiparametric factor model for electricity forward curve dynamics. *Journal of Energy Markets*, 1(3):3–16, 2008.

S. Borovkova and H. Geman. *Risk management in commodity markets: From shipping to agriculturals and energy*, chapter Forward curve modelling in commodity markets, pages 9–32. Wiley Finance Series. John Wiley & Sons Ltd, 2008.

T. Chantziara and G. Skiadopoulos. Can the dynamics of the term structure of petroleum futures be forecasted? Evidence from major markets. *Energy Economics*, 30:962–985, 2008.

L. Clewlow and C. Strickland. *Energy derivatives pricing and risk management*. Lacima Publications, 2000.

G. Cortazar and E. S. Schwartz. The valuation of commodity-contingent claims. *The Journal of Derivatives*, pages 27–39, Summer 1994.

G. Cortazar and E. S. Schwartz. Implementing a stochastic model for oil futures prices. *Energy Economics*, 25:215–238, 2003.

P. J. Dawson, A. I. Sanjuán, and B. White. Structural breaks and the relationship between barley and wheat futures prices on the London International Financial Futures Exchange. *Review of Agricultural Economics*, 28(4):585–594, 2006.

R. De Bock and J. Gijón. Will natural gas prices decouple from oil prices across the pond? Working Paper WP/11/143, International Monetary Fund, June 2011.

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

R. Engle and C. Granger. Co-integration and error correction: Representation, estimation and testing. *Econometrica*, 55(2):251–276, 1987.

R. Engle and R. Kelly. Dynamic equicorrelation. *Journal of Business & Economic Statistics*, forthcoming, 2012.

R. Engle and K. Kroner. Multivariate simultaneous generalized ARCH. *Econometric Theory*, 11:122–150, 1995.

A. Eydeland and H. Geman. Pricing power derivatives. *RISK*, September:71–73, 1998.

E. Fama and K. French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33:3–56, 1993.

J. Gabillon. The term structures of oil futures prices. Working Paper WPM 17, Oxford Institute for Energy Studies, 1991.

H. Geman. *Commodities and commodity derivatives: Modeling and pricing for agriculturals, metals and energy*. Wiley Finance Series. John Wiley & Sons Ltd, fourth edition, 2005.

H. Geman and V.-N. Nguyen. Soybean inventory and forward curve dynamics. *Management Science*, 51(7):1076–1091, 2005.

M. Grasso and M. Manera. Asymmetric error correction models for the oil-gasoline price relationship. *Energy Policy*, 35:156–177, 2007.

P. Hartley, K. Medlock, and J. Rosthal. The relationship of natural gas to oil prices. *The Energy Journal*, 29(3):47–66, 2008.

D. Heath, R. A. Jarrow, and A. Morton. Contingent claim valuation with a random evolution of interest rates. *The Review of Futures Markets*, 9(1):54–76, 1990.

S. L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6:327–343, 1993.

W. Hoeffding. Maßtabinvariante Korrelationstheorie, Schriften des Matematischen Instituts und des Instituts für angewandte Mathematik der Universität Berlin, 5, Heft 3, 179-233, 1940. [Reprinted as Scale-invariant correlation theory in The Collected Works of Wassily Hoeffding, N.I. Fisher and P.K. Sen editors, Springer-Verlag, New York, 57-107].

M. Hofert. Sampling Archimedean copulas. *Computational Statistics and Data Analysis*, 52:5163–5174, 2008.

S. Järvinen. *Essays on pricing commodity derivatives*. PhD thesis, Helsinki School of Economics, 2004.

H. Joe. *Multivariate models and dependence concepts*. Monographs on Statistics and Applied Probability. Chapman & Hall, 1997.

S. Koekebakker and F. Ollmar. Forward curve dynamics in the Nordic electricity market. *Managerial Finance*, 31(6):73–94, 2005.

L. Kullback and R. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22:79–86, 1951.

P. Kupiec. Techniques for verifying the accuracy of risk measurement models. *Journal of Derivatives*, 3:73–84, 1995.

T.-H. Lee and X. Long. Copula-based multivariate GARCH model with uncorrelated dependent errors. *Journal of Econometrics*, 150:207–218, 2009.

R. Litterman and J. Scheinkman. Common factors affecting bond returns. *The Journal of Fixed Income*, pages 54–61, June 1991.

P. Liu and K. Tang. No-arbitrage conditions for storable commodities and the modeling of futures term structures. *Journal of Banking & Finance*, 34:1675–1687, 2010.

F. Longin and B. Solnik. Extreme correlation of international equity market. *Journal of Finance*, 56:649–679, 2001.

H. Lütkepohl and M. Krätzig, editors. *Applied Time Series Econometrics*. Themes in Modern Econometrics. Cambridge University Press, 2004.

M. Manoliu and S. Tompaidis. Energy futures prices: term structure models with Kalman filter estimation. *Applied Mathematical Finance*, 9:21–43, 2002.

A. McNeil. Sampling nested Archimedean copulas. *Journal of Statistical Computation and Simulation*, 78(6):567–581, 2008.

A. McNeil and J. Nešlehová. Multivariate Archimedean copulas, d-monotone functions and l_1-norm symmetric distributions. *The Annals of Statistics*, 37(5B):3059–3097, 2009.

R. C. Merton. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(2-3):125–144, 1976.

R. B. Nelson. *An Introduction to Copulas*. Springer Series in Statistics. Springer Science+Business Media, Inc., second edition, 2006.

S. Ohana. Modeling global and local dependence in a pair of commodity forward curves with an application to the US natural gas and heating oil markets. *Energy Economics*, 32:373–388, 2010.

O. Okhrin, Y. Okhrin, and W. Schmid. On the structure and estimation of hierarchical Archimedean copulas. *Journal of Econometrics*, under revision, 2012.

D. Pilipovic. *Energy risk: Valuing and managing energy derivatives*. McGraw-Hill, second edition, 2007.

H. Reisman. Movements of the term structure of commodity futures and the pricing of commodity claims. 1991.

M. Richardson and T. Smith. A test of multivariate normality of stock returns. *Journal of Business*, 66:295–321, 1993.

P. Rousseeuw and C. Croux. Alternatives to the median absolute deviation. *Journal of the American Statistical Association*, 88(424):1273–1283, 1993.

E. Schwartz and J. E. Smith. Short-term variations and long-term dynamics in commodity prices. *Management Science*, 46(7):893–911, 2000.

E. S. Schwartz. The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance*, 52(3):923–973, 1997.

A. Sklar. Fonctions de répartition à n dimensions et leurs marges. *Publ. Inst. Stat. Univ. Paris*, 8:229–231, 1959.

C. Tolmasky and D. Hindanov. Principal components analysis for correlated curves and seasonal commodities: The case of the petroleum market. *The Journal of Futures Markets*, 22(11):1019–1035, 2002.

R. S. Tsay. *Analysis of Financial Time Series*. Wiley Series in Probability and Statistics. John Wiley & Sons, Inc., second edition, 2002.

Y. Tse and A. Tsui. A multivariate generalized autoregressive conditional heteroscedasticity model with time-varying correlations. *Journal of Business and Economic Statistics*, 20:351–362, 2002.

J. Wang and S. Zionts. Random weight generation in multiple criteria decision models, 2006. MCDM 2006, Chania, Greece, June 19-23, 2006.