Cross-Commodity Analysis and Applications to Risk Management

Reik Börger

Institute of Mathematical Finance,

Ulm University, Germany,

Email: reik.boerger@uni-ulm.de

Álvaro Cartea[*]

Commodities Finance Centre,

Birkbeck College, University of London, UK

Email: a.cartea@bbk.ac.uk

Rüdiger Kiesel

Institute of Mathematical Finance,

University of Ulm, Germany

and

Department of Statistics,

London School of Economics, UK

Email: ruediger.kiesel@uni-ulm.de

Gero Schindlmayr

EnBW Trading GmbH, Karlsruhe, Germany

December 7, 2007

1

Abstract

The understanding of joint asset return distributions is an important ingredient for managing risks of portfolios. While this is a well-discussed issue in fixed income and equity markets, it is a challenge for energy commodities. In this paper we are concerned with describing the joint return distribution of energy related commodities futures, namely power, oil, gas, coal and carbon.

The objective of the paper is threefold. First, we conduct a careful analysis of empirical returns and show how the class of multivariate generalized hyperbolic distributions performs in this context. Second, we present how risk measures can be computed for commodity portfolios based on generalized hyperbolic assumptions. And finally, we discuss the implications of our findings for risk management analyzing the exposure of power plants which represent typical energy portfolios.

Our main findings are that risk estimates based on a normal distribution in the context of energy commodities can be statistically improved using generalized hyperbolic distributions. Those distributions are flexible enough to incorporate many characteristics of commodity returns and yield more accurate risk estimates. Our analysis of the market suggests that carbon allowances can be a helpful tool for controlling the risk exposure of a typical energy portfolio representing a power plant.

1 Introduction

Managing the financial risks induced by power plants is one of the key components of the risk management of utility companies. These financial risks may be represented by certain financial futures positions, e.g. a coal-powered power plant can be thought of as a long position in electricity, a short position of coal and, due to recent efforts to regulate emissions, a short position in CO_2 emission certificates. In our analysis we will consider the interaction of all of these positions including emission certificates.

In order to evaluate the risk of these portfolios we need to analyze the dependence structure across different commodities. Based on empirical studies in Börger (2007) and Högerle (2007), we compute risk measures such as value-at-risk and expected shortfall based on different choices of multivariate distributions and apply the procedure to a given portfolio of energy commodities. As an improvement to the Normal variance-covariance approach for deducing value-at-risk in the context of commodities we propose an alternative based on generalized hyperbolic distributions that is straightforward to implement. We compare this multivariate parametric method with a copula-based approach, where we estimate the univariate distributions of the commodities under consideration first and then fit the dependence structure using copulas.

Although there are pressing needs to model power, coal, gas, oil and emission allowances simultaneously, there is very little literature on this topic. In particular, multivariate research

2

Figure 1: Q-Q plots of log-returns of futures prices: power, coal, oil, carbon (left to right)

on statistical properties of energy-related commodities is all but non-existent. Research on commodities as an asset class for investment purposes can be found in Kat and Oomen (2006) where the authors discuss the role of commodity markets as an alternative investment to bonds and stocks. They focus on different dependence measures for individual commodities and conclude that commodities, in general, seem to be statistically independent from other financial markets, but there is strong dependence within a commodity group. Further, the authors state that a multivariate Normal distribution is not appropriate to capture the dependence structure among commodities, stocks, bonds and inflation.

We will focus on the dependence structure within the group of energy-related commodities, which is of most importance to utility companies. We propose two methodologies to model and estimate the joint return distributions of various commodities and apply the results to risk management issues. Our first approach is based on explicit multivariate distributions. This requires multivariate distributions flexible enough to fit the marginal distributions of the single commodities, whilst at the same time, capturing their dependence structure. Within our second approach we fit univariate distributions first, and then model the dependence structure using copulae.

The energy commodities discussed in this article include Brent crude oil, coal, power and CO_2 emission allowances with a focus on the European market. The distributional properties of electricity prices have been examined by Cartea and Figueroa (2005) and Weron (2006) among others. Further, Eberlein and Stahl (2003) give statistical evidence for modeling electricity spot prices using generalized hyperbolic distributions and show how this distribution can be used for deriving risk capital charges based on value-at-risk. They discuss the superiority of this parametric approach to standardized alternatives that have to be applied to commodities markets in context of the German implementation of the capital adequacy directive.

From a statistical point of view, a multivariate Normal distribution is not suitable because all marginal distributions would also be Normal and the hypothesis of normally distributed returns has been rejected for assets in many markets; especially energy commodities. Fig. 1 shows the Q-Q plot of univariate log-returns of commodity futures contracts. The oil contract

in this example is close to a Normal, but for power this is not true, even less so for CO_2 prices. In fact, the statistical properties of energy commodities are very different from those of most asset classes, therefore the joint distribution is most likely described best by copulae-based methods. At the same time, univariate findings, see also Benth and Salyte-Benth (2004) indicate that distributions from the generalized hyperbolic class are well-suited to describe univariate properties, so that it is natural to consider their multivariate counterparts in addition to copula-based methods.

The rest of the paper is structured as follows. In section 2 we show how the loss distribution of a typical energy portfolio may be determined given the estimates from our two competing approaches and how to compute the corresponding risk measures. In section 3 we summarize the analyses in Börger (2007) and Högerle (2007). In section 4 we employ our results to give further directions for financial model building, pricing and hedging purposes. As a particular example, we will show how our findings can be used to evaluate the riskiness of a given energy portfolio and illustrate our methodology using portfolios that reflect the typical exposure of power plants. Finally, section 5 concludes. An appendix provides the necessary mathematical tools.

2   Portfolios and Risk measures

The value at time t, denoted by V_t, of a utility's portfolio will depend on the assets, $S_t^{(i)}$, and the number of contracts per asset, ω_i. Hence, we write the portfolio's value at time t as

$$V_t = \sum_{i=1}^{d} \omega_i S_t^{(i)}, \tag{1}$$

where $i = 1, \ldots, d$ and d represents the number of distinct assets in the portfolio. We use (1) to express the one-period loss of the portfolio as the random variable

$$X = -(V_{t+1} - V_t).$$

In order to determine risk measures we need to specify the distribution of X. This specification can be achieved by modeling the joint log-return distribution.[1] We then know the distribution of

$$x_{t+1}^{(i)} := s_{t+1}^{(i)} - s_t^{(i)}, \ s_t^{(i)} := \ln S_t^{(i)}, \ \ i = 1, \ldots, d$$

and we have

$$X = -(V_{t+1} - V_t) = -\sum \omega_i S_t^{(i)} \left(\exp(x_{t+1}^{(i)}) - 1 \right),$$

---

[1]The alternative is to model the joint return distribution $X_{t+1}^{(i)} := (S_{t+1}^{(i)} - S_t^{(i)})/S_t^{(i)}, \ \ i = 1, \ldots, d$ with $X = -(V_{t+1} - V_t) = -\sum \omega_i S_t^{(i)} X_{t+1}^{(i)}$. The analysis in Högerle (2007) and Börger (2007) find that differences between the approaches are small compared to other uncertainties regarding reliability of data and estimates.

which we can approximate for a small time horizon Δ by

$$X^\Delta = -\sum \omega_i S_t^{(i)} x_{t+1}^{(i)}. \tag{2}$$

It is clear that the approximation applied here is valid only for small log-returns.

In this paper we follow two different routes to obtain the joint log return distributions of the commodities under consideration, namely estimation of a multivariate parametric distribution and estimation of the univariate logreturn series and then combining them using a copula.

As the multivariate parametric distribution we use the Generalized Hyperbolic (GH) distribution. Recall the representation of GH distributions as Normal mean-variance mixtures, namely a random variable S is a multivariate Normal mean-variance mixture if

$$S \overset{d}{=} \mu + W\gamma + \sqrt{W}AZ \tag{3}$$

with Z a k-dimensional standard Normal, W a non-negative real-valued random variable (mixing variable) independent of Z, $A \in \mathbb{R}^{d \times k}$ and $\mu, \gamma \in \mathbb{R}^d$. If W has a generalized inverse Gaussian distribution with parameters (λ, χ, ψ), then S is GH-distributed with $\Sigma = AA^t$. Often, one can make use of the fact that S, conditional on $W = w$, is Normal with mean $\mu + w\gamma$ and variance $w\Sigma$. Further details may be found in the appendix.

One of the advantages of modeling the joint distribution of commodities with GH distributions is that we can can express, analytically, the loss distribution. Moreover, this analytical tractability allows us to study the impact of portfolio weights and try to find optimal weights (for a further discussion of this issue see appendix).

Rather than modeling the multivariate distribution directly as above, it is possible to construct a multivariate distribution by specifying marginal distributions and coupling them together with a suitable copula. The choice of copula determines the choice of dependence among the marginal distributions. An introduction to the concept of copula modeling in risk management can be found in Embrechts, Frey, and McNeil (2005). In the empirical analysis of Högerle (2007), among the many possible choices of copula function, the t-copulae was found to perform best.

Denoting by $t_{\nu,R}$ the d-dimensional t distribution with correlation matrix R and ν degrees of freedom and by $t_\nu^\leftarrow$ the quantile function of a univariate t distribution, the d-dimensional t-copula with correlation R and ν degrees of freedom is defined by

$$C_{\nu,R}^t(u_1, \dots, u_d) := t_{\nu,d}\left(t_\nu^\leftarrow(u_1), \dots, t_\nu^\leftarrow(u_d)\right).$$

Such a copula can be estimated using maximum-likelihood procedures as described Embrechts, Frey, and McNeil (2005).

Having determined the joint distribution of commodities based on marginal distributions coupled together by some well-suited copula, the portfolio distribution is not given as a

closed-form expression as above and one needs to simulate from the multivariate distribution given by the copula and compute simulated portfolio returns. This can easily be achieved by simulating a vector X according to $t_{\nu,R}$, using the fact that this distribution is a special GH distribution (with density given according to (5) in the appendix). Finally we transform the vector via $U := (t_{nu}(X_1), \ldots, t_{nu}(X_d))$ and obtain a vector U which has distribution function $C_{\nu,R}^t$ as required. The realizations of portfolio returns are computed by using Eq. (2). Thus, we approximate the true distribution by an empirical distribution based on simulation.

Having obtained the portfolio loss distribution we compute the risk measures value-at-risk (VaR) and expected shortfall (ES). Given a random variable X, representing the loss during a fixed period Δ, with distribution F_X, the value-at-risk is defined by

$$VaR_\alpha[X] = F_X^\leftarrow(\alpha)$$

for a specified confidence level α. $F_X^\leftarrow$ denotes the quantile function of F_X. The VaR is the smallest number such that the probability of the loss X exceeding the number is less than $1 - \alpha$, where α is typically above 0.9.

Expected shortfall of an absolutely continuous random variable X is defined by

$$
\begin{aligned}
ES_\alpha[X] &= \frac{1}{1-\alpha} \int_\alpha^1 VaR_u[X]du \\
&= \frac{1}{1-\alpha} \mathbb{E}\left[X \mathbb{1}(X \geq F_X^\leftarrow(\alpha))\right],
\end{aligned}
$$

where $\mathbb{1}(\cdot)$ denotes the indicator function.

We point out that whilst the VaR is the quantile of the loss distribution, the ES takes the whole tail of the loss distribution into account.

For our parametric GH distribution we can make use of the Normal mean-variance mixture representation to compute the expected shortfall. We obtain (see the appendix for a proof)

$$ES_\alpha[X] = \mu + \gamma \mathbb{E}[W] + \sigma \frac{\phi(\Phi^{-1}(\alpha))}{1-\alpha} \mathbb{E}\left[\sqrt{W}\right]. \tag{4}$$

Since we are only able to simulate the portfolio distribution in a copula model, we have to estimate the risk measures from the simulated distribution using standard estimators based on the empirical portfolio distribution. We want to emphasize that one only needs to simulate samples for the joint returns once and use these samples to deduce the portfolio distribution for different weights. This saves a considerable amount of time.

3   Statistical Analysis

3.1   Data Set

The data under consideration are daily price series of futures or forwards on power, coal, oil and carbon emission allowances. While the portfolio of an energy producing company will most likely include other financial instruments (e.g. FX positions), the energy-related commodities make up a large part of a typical portfolio of a power producing utility.

We use futures for the year 2008 with a length of 585 trading days to determine the dependence structure across commodities. The length of the data set is specified by the length of the time series that is the youngest, i.e. CO_2 certificates here. In this section we briefly discuss the specifics of each time series. The time series of power prices are taken from EEX-traded futures for the year 2008. The futures prices are quoted in EUR/MWh. Coal futures have been traded at several exchanges, among them the ICE (formerly IPE) in London and the EEX. They offer trades in coal with different points of delivery. In the following we will pick Rotterdam as an example. The futures prices are quoted in US$/t but we convert to EUR/t. While the futures market for coal has just started, there is an established market for swaps and forwards on coal.[2] An important location of trade for Brent crude oil is the ICE in London. Since there is no 2008 oil future available at the ICE, we need to construct it artificially as weighted average of monthly futures of the corresponding period. Again, we convert all prices to EUR/barrel. Emission allowances, also known as CO_2 or carbon certificates, are the youngest of the energy related commodities. They were introduced, in the context of the Kyoto-protocol, by the European Union and exchange based trading started in October 2005. One certificate allows the emission of one ton of CO_2 during a certain time period and companies need to cover their yearly emissions by certificates. The trade in allowances is divided into two periods: Period 1 is up until the end of 2007 and a Period 2 ranges from 2008 to 2012. Although allowances may be transferred within the same period, they cannot be transferred from Period 1 to Period 2. Futures contracts for allowances within one of the two periods but with different maturity dates differ only by a discount factor. In effect, this means that the 2006 and 2007 contracts for allowances in the first period are virtually the same, but different from the 2008 contract for allowances in Period 2.[3]

3.2   Summary of Statistical Results

The full statistical analysis of the copula-based approach is given in Högerle (2007) and we summarize some of them.

Analyzing the marginal distributions of power, oil, coal and carbon, we can state that for the

---

[2]Forward prices have kindly been provided by EnBW Trading.
[3]OTC prices for 2008 have been provided by EnBW Trading.

	Power	Coal	Oil	CO_2	Power	Coal	Oil	CO_2	Power	Coal	Oil	CO_2
	Data				Implied by t-copula				Implied by NIG			
Power	1.00	0.16	0.10	0.60	1.00	0.15	0.09	0.62	1.00	0.16	0.09	0.60
Coal	0.16	1.00	0.19	0.13	0.15	1.00	0.18	0.10	0.16	1.00	0.16	0.12
Oil	0.10	0.19	1.00	0.12	0.09	0.18	1.00	0.10	0.09	0.16	1.00	0.10
CO_2	0.60	0.13	0.12	1.00	0.62	0.10	0.10	1.00	0.60	0.12	0.10	1.00

Table I: Correlation estimates of year 2008 futures log-returns of power, coal, oil and CO_2 allowances for different models

univariate marginal distributions the symmetric t-, symmetric hyperbolic, NIG and skewed t-distributions give best fit to the respective data set according to likelihood-ratio-tests and model selection criteria. The marginal distributions are well-described by these distributions in line with the univariate findings in existing literature (cp. Section 1).

The t-copula has been identified as a good candidate to couple the marginal distributions to a multivariate distributions. Table I contains the correlation matrix implied by this copula compared with data estimates and can serve as one indicator of the goodness-of-fit. All correlations are well-represented with this approach. The statistical analysis for the copula method in Högerle (2007) indicates a very good description of multivariate properties of return data.

The full statistical analysis of the approach based on multivariate generalized hyperbolic distributions is given in Börger (2007) on a slightly different data set. We present updated numbers here and summarize the results.

Likelihood ratio tests hint at a multivariate Normal inverse-Gaussian distribution to describe the return data. On one hand, table I shows that implied correlations match the empirical values very precise and indicate that dependence measured by correlations is described well. On the other hand, Börger (2007) describes an issue regarding the implied marginal distributions, which tend to underestimate the tail of power and CO_2 data while overestimating the tail of coal and oil.

The averaging effect of the tail behaviour can be explained by the univariate mixing variable in the construction of GH distributions. It is the same for all marginal distributions and can only be a compromise among all commodities. While this issue might not be as prominent in equity markets, the case of commodities can be extreme. As we have shown in the introduction (cp. Fig. 1), commodity prices show returns that can range from almost Normal (e.g. oil) to heavy tails (e.g. power).

Since the copula method starts from estimating the marginal distributions, the problem is not relevant there. At the same time, correlation estimates are equally good so that the copula approach offers an advantage when fitting the data.

4 Application to Risk Management

In this section, we want to show how to apply the method to a typical energy portfolio and discuss the results. Utility companies as well as financial institutions in the commodity market need to know the riskiness of their portfolio for several reasons. Here we mention some of them:

- They can base trading strategies on the risk numbers in such a way that the financial risk is minimized.

- They can use risk numbers to set limits for traders.

- Banks have to comply to regulatory standards which are based on riskiness of portfolios.

Thus, risk numbers need to be reported on a daily or weekly basis. Of course, the time horizon for the risk number, say the value-at-risk, can be different from application to application. Trading strategies require short-term value-at-risk metrics while the regulatory standards require longer-term risk figures. We will restrict to one-day value-at-risk and one-day expected shortfall, since the statistical analysis is done on daily prices, but the mechanics are straightforward to apply to other time periods. Alternatively, scaling principles can be applied to obtain risk measures for longer holding periods. We want to mention here that the square-root scaling rule is very popular but difficulties can occur when looking at heavy-tailed data. Embrechts, Kaufmann, and Patie (2005) discuss this issue and present a method to obtain longer-period risks by scaling.

The portfolio of an energy producer depends on many variables. Key components include: size of the company, number and types of power plants, number and type of customers, hedging strategies and many more. While we could assume an arbitrary energy portfolio to analyze the company's exposure to various sources of risk, we focus on the typical building blocks and choose the two most widely used types of power plants in the industry, namely a coal-fired power plant and a gas-fired power plant. We point out that the choice of power plant is arbitrary to a certain extent, but represents a quantity that is easy to interpret. Further, it is a financial position of interest for many utility companies. Finally, banks, that do not own plants but trade in commodity markets, can interpret the portfolio as a spread contract being short electricity and long one of coal, oil or gas.

Although power plants are exposed to several types of risks, such as operational risk, volume risk and many more, in our analysis we will cover the main financial risks in the following way. First, we represent the plant by certain financial futures positions, e.g. long positions of electricity and short positions of coal. Thus, we neglect all optionalities included in the timing of production, i.e. we assume that we run the plant at electricity baseload times and prices and do not incorporate the possibility of larger earnings when producing the energy at peakload times and prices.

Second, risk metrics such as value-at-risk assume a mark-to-market valuation, i.e. compare prices of the assets today with possible prices of the assets at the end of the time period (one day in our case). This relies on the fact that the owner of the portfolio is able to sell the assets at current market prices. This need not be the case as the portfolio size can be large compared to the market and selling all the portfolio would at least influence prices. Then one would have to reduce positions step by step so that the portfolio selling is performed over several periods. Taking this into account, one has to think about an optimal closing strategy, which is not in the scope of this paper. But if the strategy is known, one would have to carry out a value-at-risk analysis for each time step as presented here.

Ignoring the limitations in the volume of trades, risk measures can be applied, strictly speaking, to portfolios whose sizes are small compared to traded volumes in the market. This would be the case, for example, when the production periods (and thus the plant's output) are small. Yet, we illustrate the procedure using typical power plants as portfolios because the economic interpretation is more intuitive than for an abstract portfolio.

The risk analysis presented below is based on multivariate distributions fitting the joint distribution of log-returns of commodity prices. The procedure of fitting and analyzing the goodness of fit has been discussed in Section 3. Similar results are obtained when working on returns instead of log-returns, but allows for exact instead of approximate value-at-risks and expected shortfalls.

4.1   Coal-fired Power Plant

A coal-fired power plant burns coal to produce electricity and as byproduct, CO_2 is emitted. Depending on the energy efficiency of the plant, only a certain fraction of the calorific value of coal is transformed into electricity. As a measure for the energy efficiency of a plant, the heat rate is often used, which is the input of heat (expressed in terms of units of fuel) needed to produce one unit (1MWh) of electricity. In the example here, we assume a heat rate of 0.33t of coal per MWh of electricity output. A byproduct of the electricity production of 1MWh is the emission of 0.9t of CO_2. Scaling this to a plant with an output capacity of 1000000MWh during a year (the portfolio can also represent a share of a larger power plant), we arrive at the following portfolio:

- Long position: 1000000MWh in power contracts. (1 contract=1MW $\hat{=}$ 8760MWh)

- Short position: 330000t in coal.

- Short position: 900000t in CO_2.

Usually, the CO_2 position is already partially covered by certificates assigned by the government. That is why we can also think of the plant as coming with additional, say, 800000t

CO_2 in certificates as a long position resulting in a net position of 100000t CO_2 short. We also consider the case when the CO_2 position is totally hedged (i.e. 0t CO_2). This latter example is comparable to the situation before the introduction of emission allowances. We find the corresponding distributions by fitting a multivariate distribution to power, coal and CO_2 prices.

Figure 2 compares the several approaches to risk estimation described above using the different portfolios.

Considering the values-at-risk and the expected shortfalls, we have the clear result that the use of the Normal distribution significantly underestimates the risks, especially for the latter, e.g. we have a shortfall of 1.22M EUR per day with the Normal distribution as compared with 1.51M EUR as can be seen in the data (without a carbon position, 98% level). The multivariate NIG also tends to underestimate the risk, but the error is much smaller, especially at the large quantiles. The copula approach gives the most conservative risk estimates. While it describes the value-at-risk precisely in all cases, it tends to overestimate the expected shortfall, in particular in the large quantiles. Yet, we have to mention, that the numbers from the copula method are estimates from simulations and bear uncertainty themselves, increasing with the level of α. Still, the use of the NIG approach is preferable in this situation.

Reducing the carbon position, i.e. closing the short position from 900000t to zero, we see a drastically increased riskiness of the portfolio. This can be explained by the strong correlation of power and carbon prices, which results in a diversification effect when taking opposite positions. Thus, for energy companies hedging their electricity price exposure it is very important to take a multi-commodity view taking also into account the exposure in emission certificates

Qualitatively, the implications of the CO_2 position is reproduced with all models considered. However, they perform differently in different situations. The method based on multivariate NIGs performs best with a large CO_2 position (lower risk), while the copula approach is preferable with no CO_2 in the portfolio (larger risk). This is due to the fact, that the rather heavy tails of power prices are reduced by the opposite CO_2 position in the portfolio. Thus, the averaging effect that takes place in the tail-estimate of multivariate GH distributions describes the tail of the portfolio well, while otherwise the heavy tails of power influence the portfolio distribution to such an extent that the deficiencies of the fit of the margins become more and more apparent, while the copula method can cope with this effect due to the precise description of marginal distributions.

Finally, we want to discuss the role of the number of CO_2 certificates. It is evident from the data that covering the CO_2 position without simultaneously reducing the electricity exposure leads to higher price risk, measured by any standard. Using the multivariate GH method, we can compute a minimum variance portfolio using relation Eq. (6) in the appendix. A short position of 481000 contracts in CO_2 would make this power plant have minimum variance.
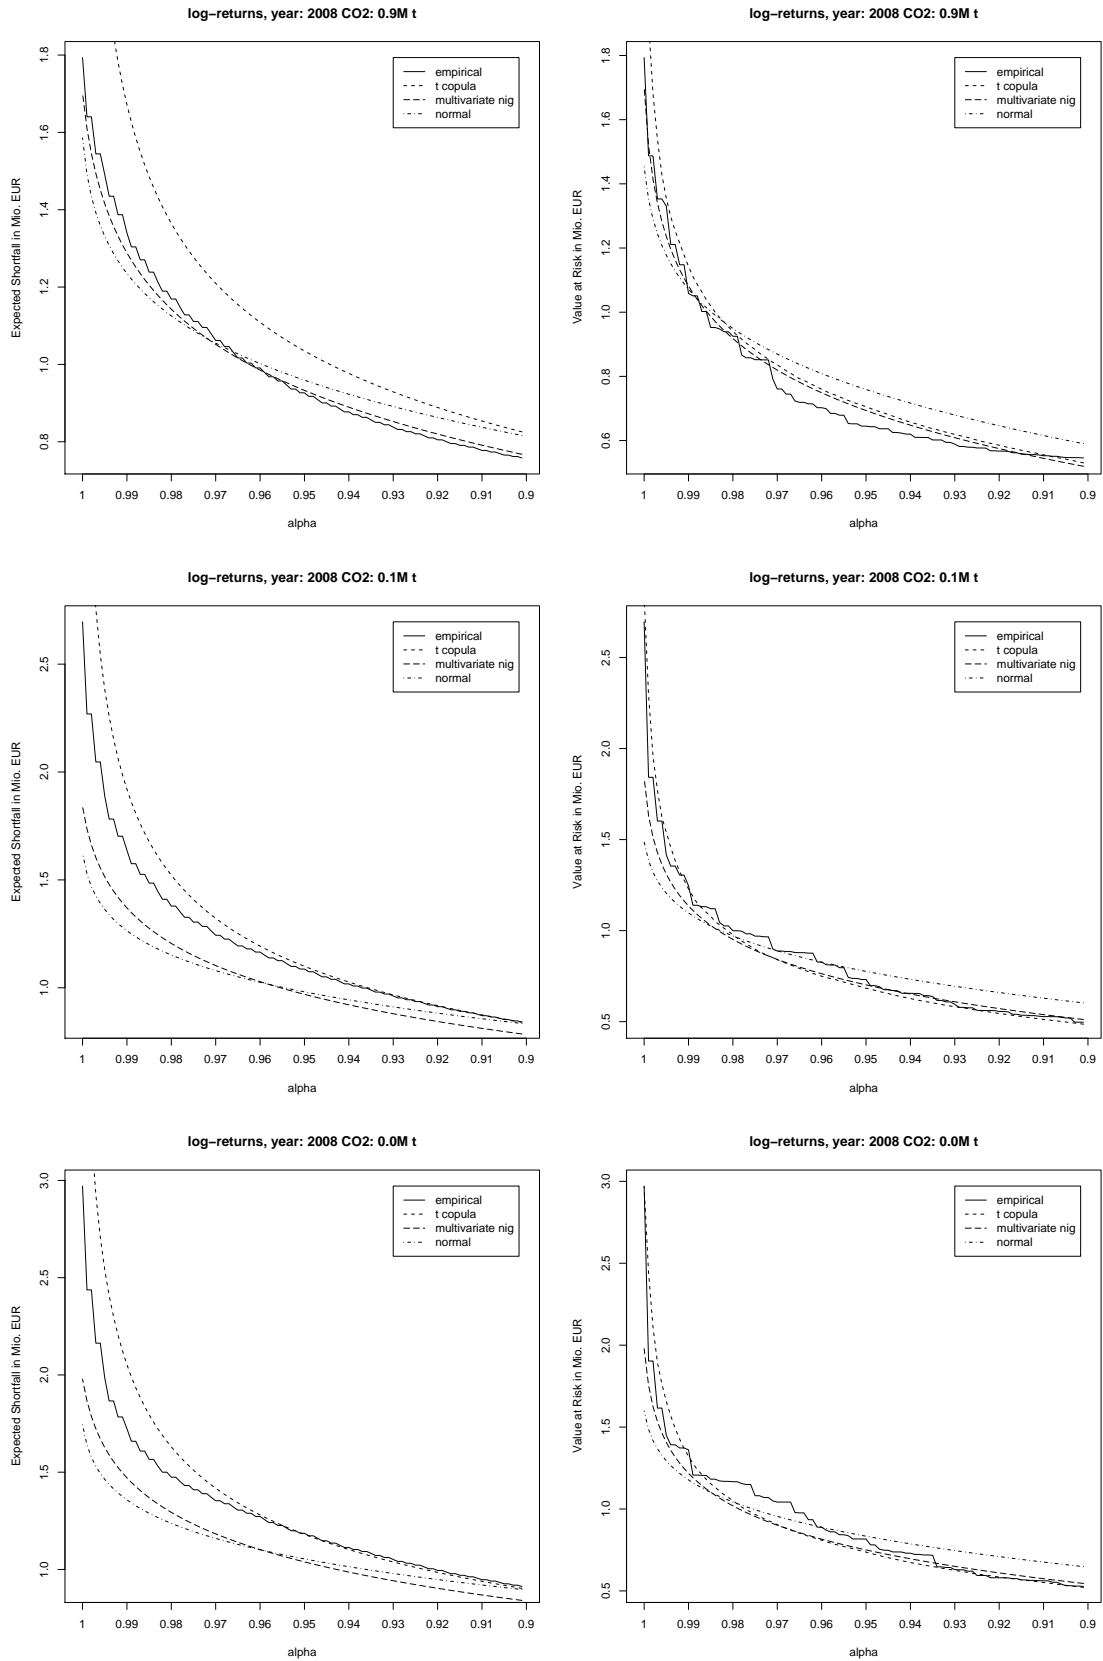
Figure 2: Expected Shortfalls (left) and Value-at-Risks (right) for coal-fired power plant portfolios

A short position of 249000 contracts in CO_2 minimizes the expected shortfall if we believe that historic returns of CO_2 prices prevail (all numbers for $\alpha = 0.95$). This means that it might be advantageous to keep an open position in emission allowances. At least it is not favorable to pursue a conservative strategy and cover all (or even more) emissions as soon as possible, but rather keep it open until the electricity exposure has been further reduced. Obviously, at some point one faces the need to buy the certificates in order to avoid the penalty by EU comission.

4.2 Gas-fired Power Plant

A gas-fired power plant burns gas to produce electricity and in the process CO_2 is released. One thermal unit of gas (1btu) has a heating value of about 0.0293 MWh. Assuming a heat rate of 68.3btu/MWh we need about 2MWh of natural gas to produce 1MWh of electricity. Burning the gas emits 0.4t of CO_2. Often, the delivery price of gas is linked deterministically to the oil price (e.g. in Germany). The exact formula is not standardized and depends also on the specific contract and the type of oil which the price is linked to. An example price formula is

$$\text{Gas in } \frac{\text{EUR}}{\text{MWh}} = \text{const} + 0.5 \cdot \text{Brent Crude Oil in } \frac{\text{EUR}}{\text{Barrel}}$$

This implies, that a short position of 2MWh of natural gas is financially equivalent to a short position of 1 Barrel of Brent Crude Oil. Scaling this to a plant that produces 1000000MWh of electricity during a year, we face the following positions:

- Long position: 1000000MWh in power contracts.

- Short position: 1000000 bbl in oil contracts.

- Short position: 400000t in CO_2 contracts.

Again, the CO_2 position is usually at least partially covered by government issued certificates, therefore we will vary the short position of certificates to 100000t and 0t. We obtain the necessary distributions by estimating the joint distribution of power, oil and CO_2 prices.

Risk numbers for such a portfolio are summarized in Figure 3. Most of the considerations apply to gas-fired power plant as to the coal-fired power plant and we will highlight differences only.

First of all, all risk values are much higher for gas than for coal though the total capacity of the plant is the same. This is largely due to the fact that at the time of analysis the oil prices have risen sharply and the total portfolio value is larger for the gas-fired plant than for the coal-fired one.

As before, the Normal distribution is clearly underestimating the risks, while - in contrast to the coal-fired plant - the multivariate NIG tends to overestimate the risk. The t-copula

gives an almost perfect match with the empirical risk values. The inapplicability of the Normal is due to the light tails, whereas the overestimation resulting from the NIG is due to the averaging effect of the tails. The average tail leads to an underestimated tail for power returns, but overestimated tail for oil returns. At the same time, oil makes up a large part of the portfolio, so that the NIG-model misspecifies the tail behaviour of the loss distribution. The t-copula on the other hand specifies marginal distributions precisely and the dependence structure as well. This explains the extraordinary good description of the risk up to very large quantiles.

The reduction of CO_2 certificates leads again to an increased risk as measured by the data and all models. The same reasoning applies as for coal-fired plants. The quality of the models changes only slightly with the varying portfolio weights. Summarizing the results, we can state that the copula model is most flexible and precise in describing the true underlying risks.

As in the previous case we compute risk minimizing portfolios. The minimum variance portfolio is established at a short position of 365000 CO_2 allowances and expected shortfalls are minimized with 362000 contracts ($\alpha = 0.95$). Again, a partially covered CO_2 position can significantly reduce the risk as compared to a closed position comparable to the time before emission trading. Such an explicit analysis is only possible in the GH case, not for the copula approach.

5 Conclusions

In this article we compared two alternative approaches for describing the joint return distribution of energy-related commodities. We showed that a copula-based approach and a multivariate generalized hyperbolic distribution (GH) are capable of describing the joint return distribution well and clearly outperform the Normal distribution.

Additionally we demonstrated how the multivariate fit of the distributions can be applied to risk management issues such as computation of risk measures. This is particularly straightforward with GH distributions due to their normality conditional on the mixture variable. The example of two different types of power plants illustrates the mechanics of the heavy-tailed, skewed distributions in connection with the dependent marginal distributions.

The most important conclusions for *risk management* are:

- GH distributions give a more realistic view on the riskiness of an energy portfolio than the Normal distribution. Since analytical formulas are available for value-at-risk and expected shortfall, the implementation is straightforward.

- A copula approach is more flexible in describing statistical properties, but the analysis of portfolios can be time-consuming.
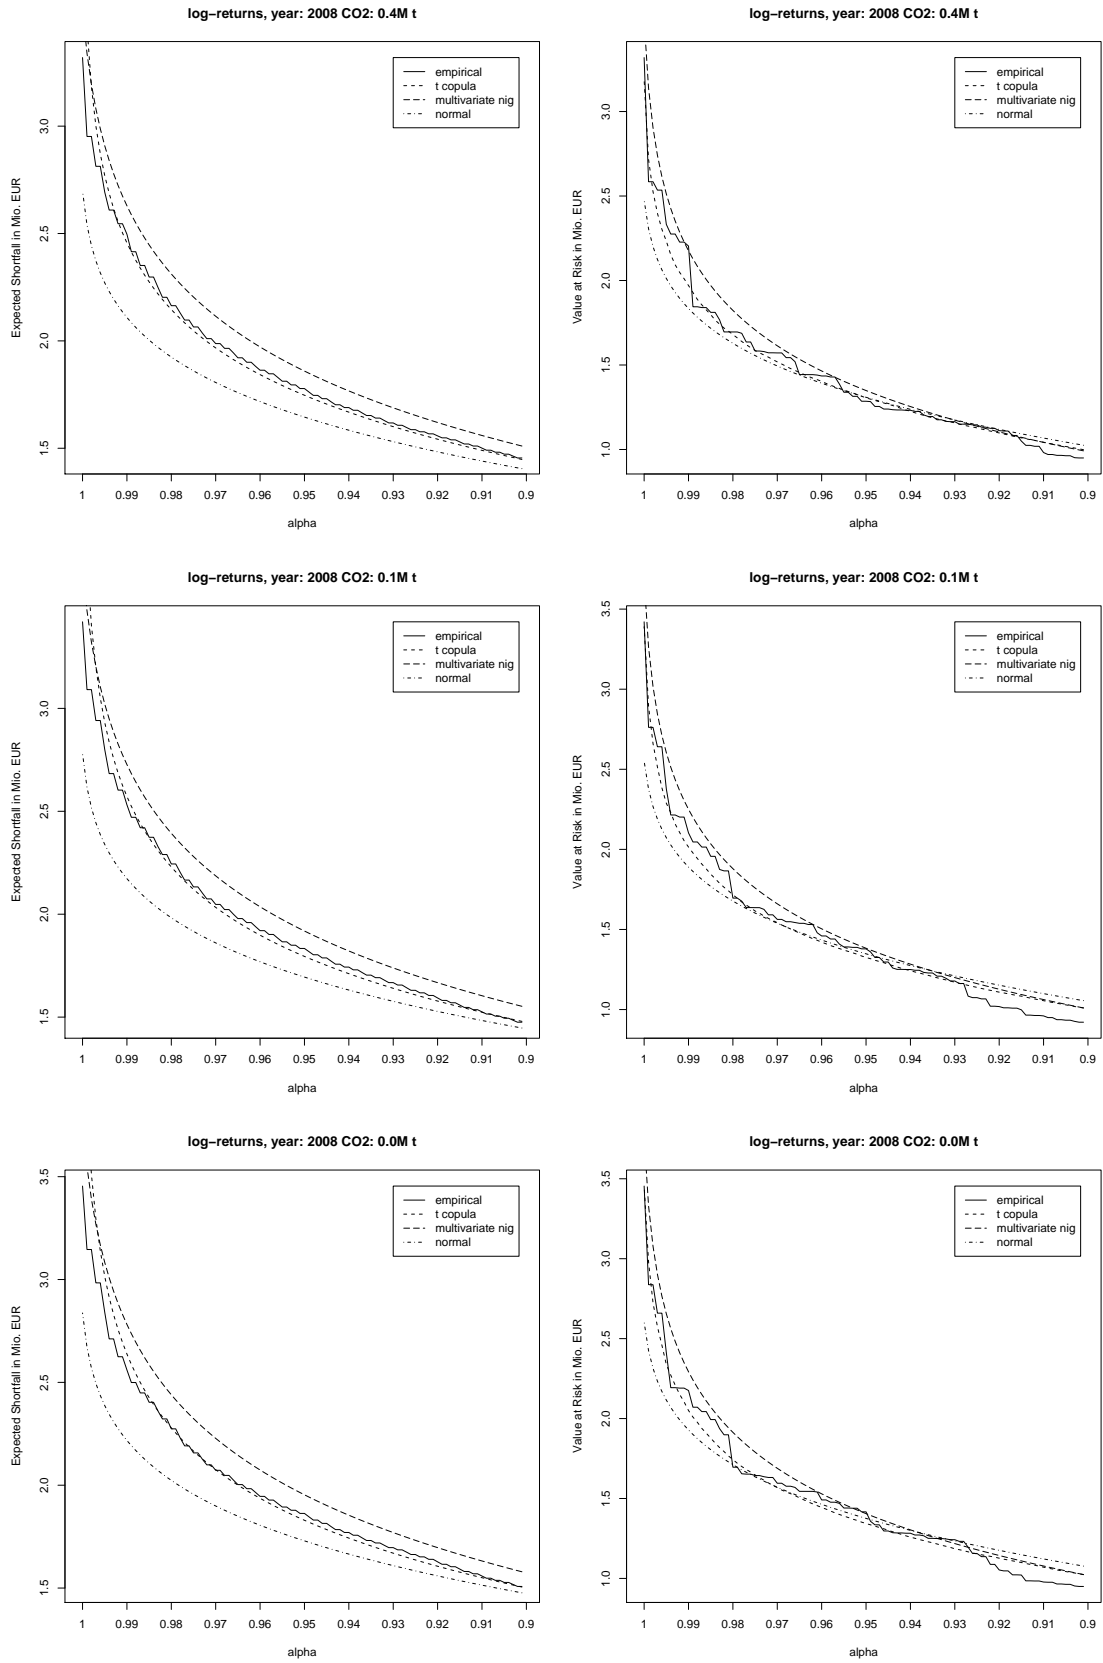
Figure 3: Expected Shortfalls (left) and Value-at-Risks (right) for gas-fired power plant portfolios

15

- The introduction of emission trading poses an additional risk factor that has to be integrated into the risk management strategy. The high correlation in the asset movements, combined with skewness, allows to reduce the financial risk for market participants by taking opposite positions in power and CO_2 certificates.

6 Appendix

We summarize some important properties of the class of GH distributions, in particular we look at the multivariate t-distribution (t), (symmetric) Normal-inverse Gaussian (NIG(-S)), and (symmetric) hyperbolic distribution (HYP (-S)). The class of GH distributions has been introduced by Barndorff-Nielsen (1978) and important properties are summarized in Barndorff-Nielsen and Blaesild (1981). Applications can be found in Embrechts, Frey, and McNeil (2005), Shiryaev (1999), Eberlein and Keller (1995), Eberlein and Prause (2000) and Eberlein and Stahl (2003).

The multivariate GH distribution is given by the joint density

$$f(x) = c\frac{K_{\lambda-(d/2)}\left(\sqrt{(\chi+(x-\mu)^t\Sigma^{-1}(x-\mu))(\psi+\gamma^t\Sigma^{-1}\gamma)}\right)e^{(x-\mu)^t\Sigma^{-1}\gamma}}{\sqrt{(\chi+(x-\mu)^t\Sigma^{-1}(x-\mu))(\psi+\gamma^t\Sigma^{-1}\gamma)}^{\frac{d}{2}-\lambda}}, \quad x \in \mathbb{R}^d \quad (5)$$

$$c = \frac{\sqrt{\chi\psi}^{-\lambda}\psi^\lambda(\psi+\gamma^t\Sigma^{-1}\gamma)^{\frac{d}{2}-\lambda}}{(2\pi)^{d/2}|\Sigma|^{1/2}K_\lambda(\sqrt{\chi\psi})}$$

with K_λ denoting the modified Bessel function of the third kind and parameters $\Sigma \in \mathbb{R}^{d\times d}$, $\mu, \gamma \in \mathbb{R}^d$, $\chi, \psi > 0$ and $\lambda \in \mathbb{R}$. Thus, a d-dimensional GH is described by $\frac{1}{2}(d(d+5)+4)$ free parameters.

Recall from Eq. (3) the representation of GH distributed random variable X as

$$X \stackrel{d}{=} \mu + W\gamma + \sqrt{W}AZ,$$

with the specifications as in 2. Observe that symmetric distributions are obtained when $\gamma = 0$. Moreover, a multivariate distribution with hyperbolic marginal distributions is given by $\lambda = 1$ and the Normal-Inverse Gaussian becomes a special case when $\lambda = -\frac{1}{2}$. The multivariate t-distribution with ν degrees of freedom can also be obtained either as a limiting case if $\gamma = 0$, $\lambda = -\frac{1}{2}\nu$, $\chi = \nu$, $\psi \to 0$ or by choosing an Inverse Gamma distribution with parameters $(\nu/2, \nu/2)$ as mixture variable (and $\gamma = 0$).

Estimation of this class of distribution can be done by Maximum-Likelihood procedures, in particular an EM-algorithm, which is an iterative scheme of maximizing conditional likelihoods. For details on the algorithm see Protassov (2004). We use the S-Plus tool provided in Embrechts, Frey, and McNeil (2005) and the freely available implementation *ghyp* for *R*.

One of the main advantages, over the copulae approach, from employing the GH approach is that we have analytical expressions for the distributions of loss of the portfolio based on log-returns or returns. In other words, we can derive the distributions of X^Δ and X, given by (2). To do so we use the fact that if the random variable Y has d-dimensional GH distribution with parameters $(\lambda, \chi, \psi, \mu, AA^t, \gamma)$ then, for a row vector $a \in \mathbb{R}^d$, aY has a one-dimensional GH distribution with parameters $(\lambda, \chi, \psi, a\mu, aAA^t a^t, a\gamma)$. Thus, the distribution of X^Δ in method one and X in method two is given by choosing $a = \left(-\omega_1 S_t^{(1)}, \ldots, -\omega_d S_t^{(d)}\right)$, and the computation of the parameters of the one-dimensional GH distribution is based on multivariate estimates.

We can compute the risk measures in the case of the alternative approach based on multivariate GH distribution. In particular, the computation of expected shortfalls turns out to be very explicit. Remember that if X has a GH distribution, it can be represented as a Normal mean-variance mixture (cp. Eq. (3), $d = 1$). We make use of the fact that, conditionally on the mixture variable $W = w$, X is normally distributed with mean $\mu + w\gamma$ and variance $w\sigma^2$. The expected shortfall can be computed by conditioning on the mixture variable W and using the result for expected shortfall in the Gaussian case, i.e.

$$ES_\alpha\left[N_{\mu,\sigma^2}\right] = \mu + \sigma \frac{\phi(\Phi^{-1}(\alpha))}{1 - \alpha}.$$

Here ϕ and Φ denote the standard Normal density and distribution functions, respectively. Further, denoting by f_W the density of the mixing variable W we compute

$$
\begin{aligned}
ES_\alpha[X] &= \frac{1}{1 - \alpha} \mathbb{E}\left[X \mathbb{1}(X \geq F_X^\leftarrow(\alpha))\right] \\
&= \int_\mathbb{R} \frac{1}{1 - \alpha} \mathbb{E}\left[X \mathbb{1}(X \geq F_X^\leftarrow(\alpha)) | W = w\right] f_W(w) dw \\
&= \int_\mathbb{R} ES_\alpha\left[N_{\mu+w\gamma, w\sigma^2}\right] f_W(w) dw \\
&= \int_\mathbb{R} \left(\mu + w\gamma + \sqrt{w}\sigma \frac{\phi(\Phi^{-1}(\alpha))}{1 - \alpha}\right) f_W(w) dw \\
&= \mu + \gamma \mathbb{E}[W] + \sigma \frac{\phi(\Phi^{-1}(\alpha))}{1 - \alpha} \mathbb{E}\left[\sqrt{W}\right].
\end{aligned}
$$

We discuss three special cases:

- Normal: If $\gamma = 0$ and $W = 1$, X is Normal with mean μ and variance σ^2 and the expected shortfall has the well-known form

$$ES_\alpha[X] = \mu + \sigma \frac{\phi(\Phi^{-1}(\alpha))}{1 - \alpha}.$$

- t: If $\gamma = 0$ and W is inverse gamma distributed with parameters $\frac{\nu}{2}$ and $\frac{\nu}{2}$, X is t-distributed with ν degrees of freedom. Here, we have

$$\mathbb{E}[W] = \frac{\nu}{\nu - 2}, \quad \mathbb{E}\left[\sqrt{W}\right] = \sqrt{2\nu} \frac{\Gamma\left(\frac{\nu-1}{2}\right)}{\Gamma(\nu/2)}.$$

17

- If W has generalized Inverse Gaussian distribution with parameters λ, χ, ψ, we obtain the class of GH distributions with NIG ($\lambda = -\frac{1}{2}$), HYP ($\lambda = 1$) and symmetric distributions ($\gamma = 0$) as special cases. Here we have

$$\mathbb{E}[W] =$$

Barndorff-Nielsen, Ole E., and P Blaesild, 1981, Hyperbolic distributions and ramifications: Contributions to theory and application, *Statistical Distributions in Scientific Work* 4, 19–44.

Benth, Fred Espen, and Jurate Salyte-Benth, 2004, The normal inverse Gaussian distribution and spot price modelling in energy markets, *International Journal of Theoretical and Applied Finance* 7, 177–192 .

Börger, Reik, 2007, Energy-Related Commodity Futures: Statistics, Models and Derivatives, Ph.D. thesis.

Burger, Markus, Bernhard Klar, Alfred Müller, and Gero Schindlmayr, 2004, A spot market model for pricing derivatives in electricity markets, *Quantitative Finance* 4, 109–122.

Cartea, Álvaro, and Marcelo G. Figueroa, 2005, Pricing in Electricity Markets: a Mean Reverting Jump Diffusion Model with Seasonality, *Applied Mathematical Finance* 12, 313–335.

Eberlein, E., and U. Keller, 1995, Hyperbolic distributions in finance, *Bernoulli* 1, 281–299.

Eberlein, E., and K. Prause, 2000, The generalized hyperbolic model: Financial derivatives and risk measures, Working paper, Mathematical Finance-Bachelier Congress.

Eberlein, Ernst, and Gerhard Stahl, 2003, Both sides of the fence: a statistical and regulatory view of electricity risk, *Energy and Power Risk Management.*

Embrechts, Paul, Rüdiger Frey, and J. Alexander McNeil, 2005, *Quantitative Risk Management - Concepts, Techniques, Tools.* (Princeton University Press).

Embrechts, Paul, Roger Kaufmann, and Pierre Patie, 2005, Strategic Long-Term Financial Risks: Single Risk Factors, *Computational Optimization and Applications* 32, 61–90.

Högerle, Wolfgang, 2007, Multivariate Commodity Modelling, Master's thesis Ulm University.

Kat, Harry M., and Roel C.A. Oomen, 2006, What Every Investor Should Know About Commodities, Part I: Univariate Return Analysis, Working paper, The alternative investment research centre.

Protassov, R. S., 2004, EM-based maximum likelihood parameter estimation of multivariate generalized hyperbolic distributions with fixed λ, *Statistics and Computing* 14, 67–77.

Schindlmayr, Gero, 2005, Modelling Electricity Spot Prices - A Regime-Switching Approach, *10th Symposium on Finance, Banking, and Insurance, Karlsruhe.*

Shiryaev, Albert N., 1999, *Essentials of Stochastic Finance: Facts, Models, Theory.* (World Scientific Publishing Company) 1st edn.

Weron, Rafal, 2006, *Modelling and forecasting electricity loads and prices: A statistical approach.* (Wiley & Sons).